

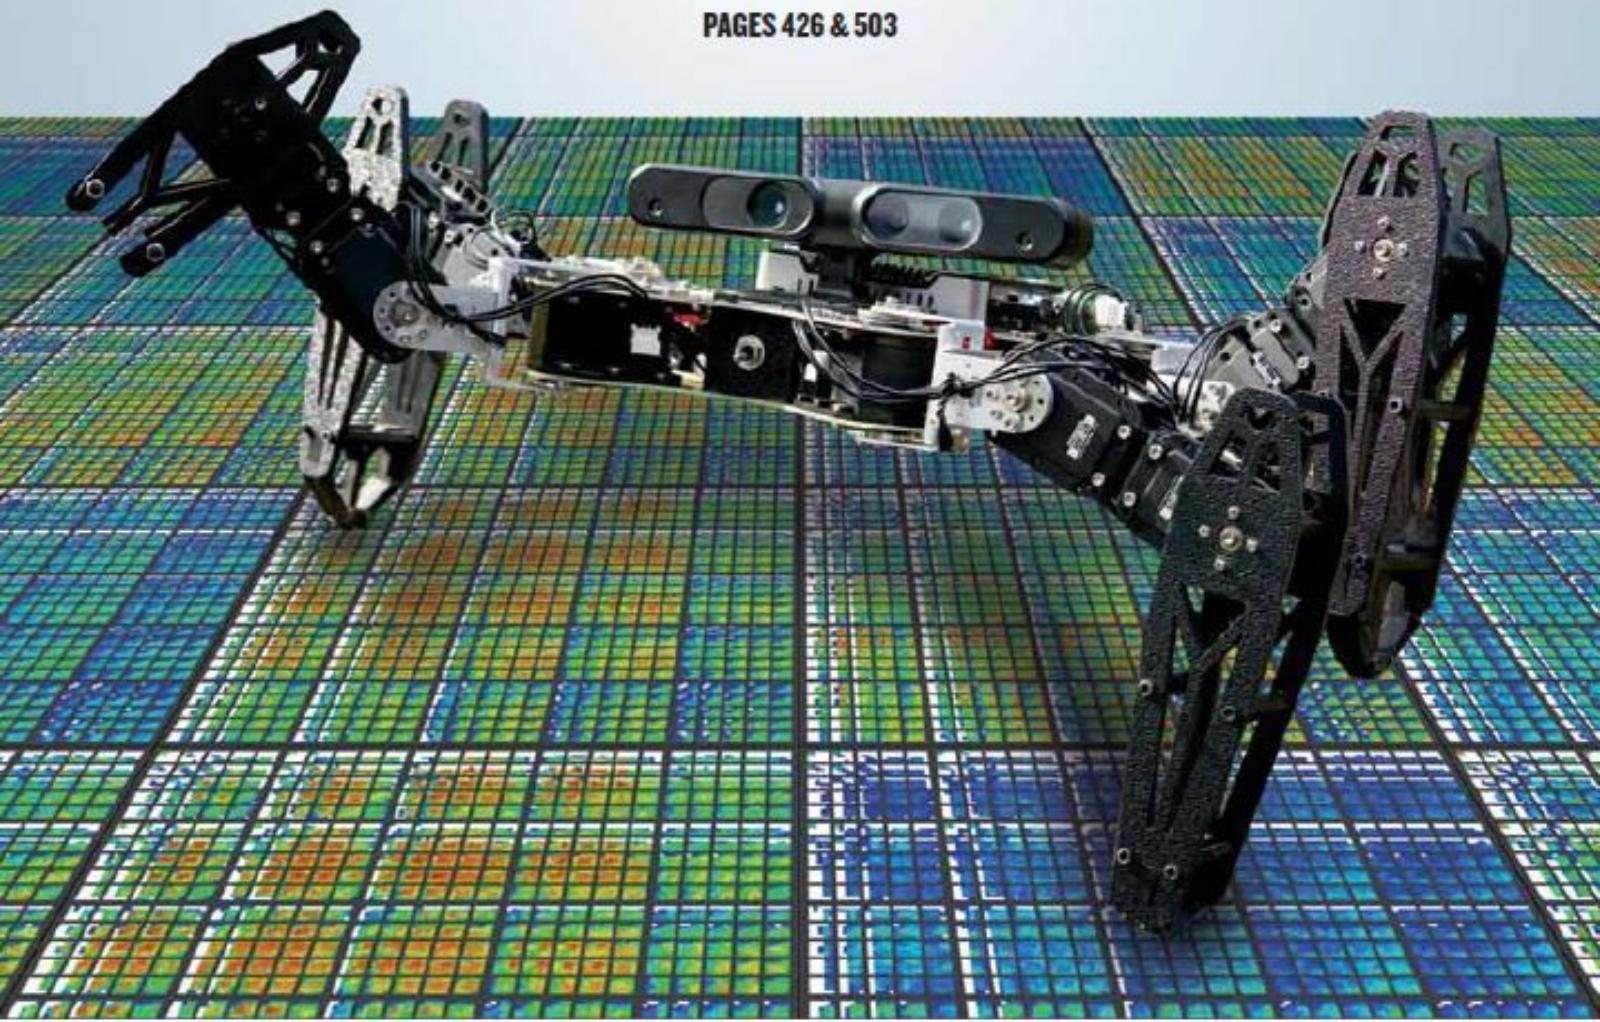
nature

THE INTERNATIONAL WEEKLY JOURNAL OF SCIENCE

Back on its feet

Using an intelligent trial-and-error learning algorithm this robot adapts to injury in minutes

PAGES 426 & 503



COGNITION

**WHY FISH NEED
TO BE CLEVER**
*Social behaviours need
plenty of brainpower*

PAGE 412

ARTIFICIAL INTELLIGENCE

**LIVING WITH
ROBOTS**
*AI researchers' ethics
prescriptions*

PAGE 415

HUMAN EVOLUTION

**ANOTHER FACE
IN THE CROWD**
*A new hominin from
Ethiopia's middle Pliocene*

PAGES 432 & 483

NATURE.COM/NATURE

28 May 2015 £10

Vol. 521, No. 7553



9 770028 083095

THIS WEEK



EDITORIALS

WAKING UP The scientific papers that come to life after decades p.394

WORLD VIEW Eating insects must be worthwhile not just useful p.395

CLEVER CROWS Crafty birds look after and reuse their tools p.397

Trading places

Scientists have a valuable part to play in clarifying the impacts of a proposed trade treaty between the United States and Europe.

The world's two biggest economies, the United States and the European Union (EU), are negotiating an accord — the Transatlantic Trade and Investment Partnership (TTIP). They argue that it could boost the world economy by more than US\$300 billion, and create millions of jobs. But at what price?

Free-trade agreements have historically boosted economic growth by eliminating border tariffs and opening up markets. But TTIP is different. Few significant tariff barriers exist between the United States and the EU, so the proposed accord focuses on reducing the economic impact of 'non-tariff barriers' — or, in plain language, reforms of standards and regulations on everything from the environment and public health to agriculture and pharmaceuticals. These could have wide-ranging and profound impacts, for good and bad.

There are clear benefits to be had from greater harmonization between the United States and the EU in some regulatory matters. Streamlining pharmaceutical regulation, for instance — so that an EU-approved product could be sold in the United States and vice versa — would reduce red tape and duplication. It would also boost research cooperation, cutting the costs of developing drugs and, ultimately, lowering prices. But other provisions of the accord could give drug companies more say in pricing and reimbursement policies — and so might result in higher medical costs and reduced access to health care. The humanitarian organization Médecins Sans Frontières (Doctors Without Borders) has also expressed concern that the impact of free-trade agreements on intellectual property increasingly threatens the access of millions of people in poor countries to affordable medicines.

As discussed on page 401, a major concern for scientific policy is that greater harmonization of EU and US policies might result in an overall decrease in regulatory standards on the environment, food safety or data protection, for example. Some say that such fears are overblown. Assessing what might be the true benefits and risks is not easy. Secrecy surrounds the negotiations, and the few details that are available usually come from unauthorized leaks. This is no way to conduct debate on such important matters. Civil society, including scientists and scientific organizations, must continue to badger for more access and information. At the very least, society should be given a similar level of access to that already granted to industry groups.

EARTH INCORPORATED

Beyond individual provisions, the nature of the proposed accord itself raises fundamental questions. Some leading economists — including Jeffrey Sachs at Columbia University in New York, previously a strong proponent of globalization, and Joseph Stiglitz, a 2001 economics Nobel laureate, also at Columbia — have called for TTIP to be rejected. They argue that it is not a trade agreement at all, but an undemocratic way for corporate lobbies to impose a narrow form of globalization — one that is focused more on commercial interests than on creating global systems to address such challenges as environmental and health crises (including

biodiversity loss) and reducing economic inequality.

One proposed provision — a mechanism for settling disputes that is designed to protect investors — deserves much wider debate. It would allow companies to bring lawsuits against sovereign states in private arbitration courts to obtain financial compensation for any regulatory actions that could harm their 'expected' profits. Similar mechanisms have already allowed the tobacco company Philip Morris to sue Uruguay

"Science policy-setting is not, and must not be, chiefly about trade."

over the introduction of national anti-smoking measures. And Swedish energy company Vattenfall sued the German government in 2009 for introducing stricter environmental controls on coal-fired power plants; that case was settled out of court in 2010. The proposals mean that France, for example, might open

itself up to lawsuits over its ban on fracking or its phasing out of a class of pesticides thought to harm bees and other insects.

The existence of such a legal mechanism at the highest levels of US and EU politics could have a dangerous and chilling effect, making countries think twice before introducing stricter regulations to protect public health or the environment, for example. It offers a ready-made route for corporate interests to usurp the right of sovereign states to act in the best interests of their citizens. In February, the European University Association expressed concern that TTIP could harm the ability of authorities to determine the shape of their higher-education systems.

The use of sound science to set regulations that affect trade is to be encouraged. But the science is not always unequivocal, and it must by no means be the only consideration. The practices of individual nations are forged from their own history and culture, resulting in different approaches to how they structure health care, agriculture, food or environmental systems — and in how these are shaped by government and the market, and to what extent. National attitudes to science and technology are formed in a similar way; for example, in the level of risk people are willing to accept, or the ethical limits that such attitudes place on research or medical practices.

That is why regulations differ worldwide. The dominant globalization ideology too often sees cultural differences, and citizens' rights to determine the sort of society they desire, as trade barriers that must be overcome. This is misguided. Science policy-setting is not, and must not be, chiefly about trade.

Done right, the harmonization of regulations and standards between the United States and the EU could bring enormous benefits. But the potential risks it carries for increased deregulation, and for the role and diversity of science-related policies in democratic societies, deserves much greater scrutiny — and much more transparent debate. Scientists, social scientists and their representatives, such as national academies, could perform a public service by contributing to and helping to clarify the many detailed scientific, technical, regulatory and social aspects of this complex and wide-ranging accord. ■

Wakey wakey

Sleeping-beauty papers offer hope that authors of uncited works are in good company.

How many of the research papers published in this week's *Nature* — or indeed any scientific journal — will go on to change the world? How many will at least make a sizable dent in their academic field, inspire future work and perhaps overturn what has gone before? Citations by other researchers are the currency of modern science, a mark of professional approval that indicates influence. But how long should one be expected to wait for them? Conventional wisdom says that the reach of a publication can be gauged by how many citations it attracts in the first five years. Which gives this week's authors until 2020 or so.

Vincent Van Gogh had no time for conventional wisdom. The artist was famously ignored in his lifetime, yet his work *L'Allee des Alyscamps* sold for US\$66 million in New York earlier this month. The books of Herman Melville were out of print when he died, and 1851's *Moby-Dick* did not surface in the public consciousness until years later. Art and science are not so different. Gregor Mendel was not recognized as the father of modern genetics until decades after his experiments with peas.

Such cases sound extreme, but a study offers the presently unloved hope that delayed recognition might not be so rare after all. After analysing a database of 22 million academic papers, researchers have identified plenty of works that went for decades before they were recognized and cited as important (Q. Ke *et al.* *Proc. Natl Acad. Sci. USA* <http://dx.doi.org/10.1073/pnas.1424329112>; 2015). One paper — entitled 'Concerning adsorption in solutions' — lay undisturbed for almost a century. Its author, the German chemist Herbert Freundlich, published the work in 1906 and died in 1941. His citation spike came in 2002.

The apparent snub did little to damage his career — Freundlich went on to be widely recognized (and cited) as a pioneer of colloids research. And an initial lack of citations for a 1935 paper on quantum mechanics did not reflect the standing of its authors: physicists Boris Podolsky, Nathan Rosen and one Albert Einstein. Only in 1994 did this publication start to be cited extensively.

Scientific papers typically accrue citations steadily, peak and then decline. Those that at first lie dormant, before being discovered and enjoying a late surge, are dubbed sleeping beauties. In many cases, the awakening comes when the published research finds applications in a different field, such as when statistical methods acquire a use in biology. Some papers were ahead of their time, and described techniques that could not be exploited properly until the creation and curation of large modern databases.

"Plenty of works went for decades before they were recognized and cited as important."

But creators of sleeping papers be warned: there is no guarantee that your prince will come. Although the latest study indicates that late-flowering papers are more common than previously thought, there remain plenty that are never cited — and never will be.

Filippo Radicchi, a researcher in complex networks at Indiana University Bloomington who worked on the study, says: "I expect, if you look at a paper that is 10 years old [and not cited] my guess is it will continue to have zero citations forever" (see *Nature* <http://doi.org/4tb>; 2015). He is now trying to identify the papers that wake the dormant studies from their slumber with an important citation.

Citation analysis is an increasing fact of academic life, and this study demonstrates, yet again, that the bare figures do not — and can never — show the full story. Some impact, and personal achievement, is simply difficult to measure, even during a productive career.

Herbert Freundlich achieved much in life, and more after his death. But he did not get everything he wanted. As his obituary in *Nature* noted, he was a talented but frustrated musician. "He abandoned music for chemistry," it said, "when he concluded that he would never be a great composer". There is still time, Herbert. ■

Silicon smarts

A package of articles in Nature assesses the state of artificial-intelligence research.

When a select band of computer scientists met at Dartmouth College in Hanover, New Hampshire, in 1956 to begin work on a field they called 'artificial intelligence', they were optimistic, to say the least. Their founding principle of developing machine intelligence was based on an assumption that human intelligence could itself be well characterized. They argued that: "Every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it."

Ask ten people to define human intelligence and you will get at least eleven answers. To a philosopher, intelligence is the absence of a lack of intelligence. To psychologists it is what intelligence tests measure.

Yet despite this fuzziness, the nature of artificial intelligence, in popular culture at least, is sharply defined: computers and robots that can think and act like a human, and that have the potential to outthink and counteract us in most situations. That is probably why many people are disappointed with what even the most advanced robots can achieve, certainly compared with the impressive abilities of even the youngest humans. In their minds, Mozart was composing and performing music at five years old whereas robots can barely fold a towel. The pre-eminence of humankind, it seems, is assured.

And yet, break down the holistic expectation of intelligence into a series of distinct (if overlapping) abilities, and the machines fare somewhat better. In a research paper on page 503, scientists define intelligence as the ability to predict the future. And they have built machines that can do it pretty well. Or at least they have built robots that can analyse the past to plan how to modify their own future behaviours if they are to continue functioning. The work's implications for the continuing survival of feeble humanity are described in a News & Views article on page 426.

Continuing the theme, a series of Comment articles starting on page 415 assesses the current state of debate over how society should respond, regulate and interact with intelligent machines. From autonomous weapons, which could be 'clever' enough to distinguish friend from foe and act accordingly, to medical diagnoses based on rapid and accurate analysis and interpretation of health-care data, these machines may not yet be classed as fully intelligent, but they are reaching a point at which they can mimic and potentially outperform specific 'intelligent' human abilities. What should be done? In the case of drones and other armed intelligent machines, decision time is looming.

Finally, a string of Review articles make up a *Nature* Insight on machine intelligence, starting on page 435. From machine-learning techniques and evolutionary computation to the design and construction

of malleable robots inspired by nature, the selection offers both a primer to the uninitiated and a useful summary of the state of the art. It is all, of course, essential reading. The machines, after all, are getting smarter. We should keep up. ■



Eat insects for fun, not to help the environment

Insects are an excellent source of sustainable protein, but people will only be persuaded to eat them if they seem appealing, says Ophelia Deroy.

As the world searches for a more sustainable future for its growing population, there is increasing interest in getting more people to eat insects. Earlier this month, former United Nations secretary-general Kofi Annan told *The Guardian* newspaper that “eating insects is good for the environment and balanced diets”. This backed the view of a widely cited 2010 report from the UN Food and Agriculture Organization that stressed insects’ “exceptional nutritional benefits” and “fewer negative environmental impacts” when compared with “many mainstream foods” (see go.nature.com/6ln9dw).

So far, there is little sign that these messages are being heeded. Turning Westerners into insectivores has joined the long list of challenges that require behaviour change. As such, it holds broader lessons for other attempts to convert people to more sustainable lifestyles, and is a useful case study.

Many people already eat insects. There is evidence that insects have been a continuous part of our diet since the early hominins, and they are still eaten widely in south and east Asia, Africa, and South and Central America. (Although the practice is declining in some places as people switch to ‘aspirational’ Western lifestyles.)

Western policy-makers and the media often send messages that lean heavily on this argument: they do it, so why can’t we? These messages use rational reasons to try to overturn the presumed major psychological objection to eating insects — disgust. This strategy assumes that the revulsion people feel when presented with, say, house-fly pupae (protein content 62%), is a cognitive process that can be addressed through education.

There is little evidence that this strategy is working. Worse, appealing to reason and responsibility reinforces a dilemma in the minds of consumers: many know that they can, in principle, eat insects, and perhaps that they should, but very few are willing to do so.

What if it is not disgust that stops people eating insects? Disgust is linked to contamination and fear of disease. Insects are considered disgusting to eat, the theory says, because insects themselves eat ‘dirty’ foods. Yet many Westerners are happy to consume lobster, which scavenge from the sea floor, and pigs, which eat slops. Many insects, including some grasshoppers and ants, have the same diet as sheep.

We should think less about combating disgust and more about appealing to taste. Most of the insects eaten in the world are cooked as part of interesting preparations that make them a genuine competitor to other foods, and often a more attractive option. These insects are eaten by choice, not necessity. This obvious fact is missed by most of the current research and policies.

Taste is affected by more than the flavour and smell of food. Also important is colour, the other visual images associated with the food and the name it is presented under. The re-naming of the (rather ugly) Patagonian toothfish as Chilean sea bass, for example, led to a sharp increase in sales. In one of the few studies to have been conducted so far, Belgian consumers were shown to accept insects (mealworms and house crickets) more readily when they were prepared using familiar flavours (R. Caparros Megido *J. Sens. Stud.* **29**, 14–20; 2014).

Chefs and others who work intimately with food know the importance of perception. I would never have eaten crickets if it was not for the beautiful golden dust that the chef sprayed them with before placing them on the side of a salad. It made them look unusual, and rare, but also made me anticipate a firm crunchy texture and a sharp taste. These expectations made the experience of eating them much less scary, and really enjoyable. And I do not remember it as ‘doing something good for the environment’ or ‘eating an insect alternative to animal protein’.

If we are serious about broadening the appeal of insects as food — and we should be — then the images we present to consumers should not be of industrially farmed meat substitute.

Telling citizens and consumers that the food industry will have licence to add insect matter to food products will make them worried and sceptical. What is insect matter after all — which insects? How were they grown? Food is a question of trust. Insects that are farmed industrially, or imported from long distances by road, air or

sea, might not be such a great help for the environment. We should focus instead on local insects. Studies suggest that people will be more likely to try a bee ice cream from their local honey producer than a stick of grilled scorpion from far away.

We must stress the importance of cooking and recipes. And we should even move on from using the too-broad term insects. We do not eat mammals and birds; we eat cows, sheep and chickens. Insects could be identified by their real names, such as house crickets and wax-moth larvae. And we should not forget those with a sweet tooth: many insects lend themselves naturally to desserts.

Most importantly, before we try to change the minds of consumers, we must understand their objections. And to overcome these objections, food scientists, chefs and psychologists must work together to make insect dishes appeal as food, not as a way to save the planet. ■

THE IMAGE WE PRESENT TO **CONSUMERS** SHOULD NOT BE OF INDUSTRIALLY FARMED **MEAT** **SUBSTITUTE.**

NATURE.COM
Discuss this article
online at:
go.nature.com/r7zcot

Ophelia Deroy is a researcher in the Centre for the Study of the Senses at the School of Advanced Study, University of London.
e-mail: ophelia.deroy@sas.ac.uk

CLIMATE SCIENCE

Less sea ice, more Siberian snow

Shrinking Arctic sea ice leads to heavier snows in western Siberia. This is some of the first evidence for how low sea-ice levels in the Arctic autumn affect precipitation in neighbouring regions the following winter.

A team led by Martin Wegmann at the University of Bern, Switzerland, looked at snow measurements from 820 locations across Russia, and modelled how moisture flows between the ocean and the atmosphere. They found that the Barents and Kara seas, north of the Ural Mountains, released more moisture into the air when the seas were not covered with ice. That moisture then ends up as snow in northern Russia.

If sea ice continues to dwindle as expected, Eurasia and other continents may need to brace for heavy winter snowfalls.

Environ. Res. Lett. 10, 054015 (2015)

ANIMAL BEHAVIOUR

Crafty crows keep their tools handy

A crow species renowned for its use of tools has a time-saving trick — stashing the same tool for future use.

New Caledonian crows (*Corvus monedulaoides*) invest much time and energy turning sticks into hooks to extract food from small cavities. To see what happened to the tools when they were not in use, Barbara Klump and Christian Rutz at the University of St Andrews, UK, and their colleagues offered crows food hidden in holes in a block of wood. They found that, while they ate, the crows usually pinned their tool underfoot (pictured) or stashed it in a hole between uses. Crows foraging up high — where a fallen tool would be harder to retrieve — were more apt to store tools in holes than



were birds lower down.

Such safekeeping probably mitigates the cost of tool-making, allowing this complex behaviour to evolve, the authors say.

Proc. R. Soc. B 282, 20150278 (2015)

CANCER GENETICS

Skin riddled with cancer mutations

More than one-quarter of cells in ageing, sun-exposed skin carry mutations that are known to drive cancer — even though the skin continues to function normally.

To understand how healthy cells can mutate to form cancers, Philip Jones of the University of Cambridge, UK, Peter Campbell of the Wellcome Trust Sanger Institute in Hinxton, UK, and their colleagues sequenced 74 cancer-linked genes in biopsies taken from excess eyelid skin removed from four healthy people aged over 55. More than 25% of the normal cells each contained thousands of cancer-causing mutations caused by sunlight — a similar frequency to that seen in many skin-cancer cells. These cells might form a reservoir with the potential to transform into tumours.

Cancer drugs that target these mutations could damage normal cells with the same mutations, the authors suggest.

Science 348, 880–886 (2015)

METABOLISM

Potential obesity drug from vine

A compound found in the roots of the 'thunder god' vine could be a weight-loss drug, a study in mice suggests.

Umut Ozcan of Boston Children's Hospital in Massachusetts and his colleagues used a database

SOCIAL SELECTION

Popular topics on social media

Genomics paper probed on Twitter

A recent Twitter conversation casting doubt on the conclusions of a genomics study has revived a debate about how best to publicly discuss possible errors in research. Yoav Gilad, a geneticist at the University of Chicago in Illinois, last month wrote on Twitter that fundamental errors in the design and data analysis of a December 2014 study led to an unfounded conclusion about the genetic similarities between mice and humans. Gilad and his co-author Orna Mizrahi-Man, a bioinformatics researcher also at the University of Chicago, have since detailed their data reanalysis in the open-access journal *F1000Research* (in which articles are openly peer-reviewed after publication).

Michael Snyder, a geneticist at Stanford University in California and co-author of the original paper, stands by his team's study and its conclusions.

F1000Research 4, 121 (2015)

of gene expression in human cells treated with various molecules to look for compounds that could increase the brain's sensitivity to leptin, a hormone that suppresses appetite. The screen fished out a compound called celastrol, found naturally in the roots of thunder god vine (*Tripterygium wilfordii*). Its skinned roots are used in traditional Chinese medicine to treat inflammation.

The team found that obese mice treated with celastrol ate less than untreated obese mice, and they lost up to 45% of their body weight. Genetically engineered animals that lacked normal leptin responses did not lose weight, however. Treatment of obese mice with celastrol also activated leptin signalling and reduced blood-sugar levels, with no obvious signs of toxicity.

Cell <http://doi.org/4sm> (2015)

MATERIALS

Printed spider webs get tough

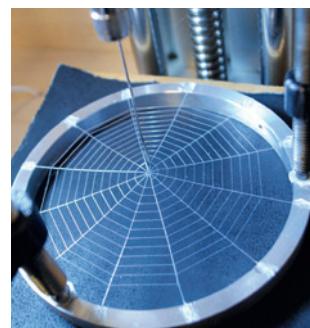
Using a 3D printer, researchers have created spider-web analogues out of elastic polymer threads and tweaked their architectures to

maximize the webs' strength.

Markus Buehler at the Massachusetts Institute of Technology in Cambridge and his colleagues printed synthetic webs (pictured) and added various sizes of load to test the webs' strength. The authors found that they could strengthen the web by adjusting the diameter of the threads radiating out from the middle, and that of the threads that spiral around the web. They did so while keeping the mass and general geometry of the web constant.

Strong synthetic webs could be used in applications such as reinforcing industrial materials, the authors say.

Nature Commun. 6, 7038 (2015)



NATURE.COM

For the latest research published by *Nature* visit: www.nature.com/latestresearch

SEVEN DAYS

The news in brief

POLICY

Animal drug trail

Drug companies that sell antibiotics to US livestock farmers may be required to track where their drugs are going. Under a new rule proposed by the US Food and Drug Administration on 19 May, drugmakers would have to collect data on how their drugs are being sold and used in different livestock. The agency plans to use the data to study how antibiotic usage in animals leads to the development of antibiotic-resistant bacteria. It is accepting comments on the rule until August.

EVENTS

LHC breaks record

The Large Hadron Collider (LHC) smashed protons together at a record energy of 13 teraelectronvolts on 20 May. After a two-year hiatus, researchers at CERN, Europe's particle-physics laboratory near Geneva, Switzerland, are gearing up to start taking data at this energy in early June. The latest collisions were designed to calibrate the systems that protect the machine and detectors from stray particles. Although the LHC's detectors were not fully operational during the tests, experiments including ALICE captured the first bursts of particles produced by collisions at this energy.

Mystery mission

The US Air Force launched its uncrewed spaceplane, the X-37B, from Cape Canaveral, Florida, on 20 May for a secretive classified test. It is the fourth flight of the vehicle, which can remain in low-Earth orbit for more than a year — much longer than the space shuttle could. The vehicle's research purposes



DAVID MCNEW/GETTY

California deals with massive oil spill

An oil spill from a broken pipeline near Refugio State Beach in California, which has contaminated 14 kilometres of coastline, led the state's governor to declare a state of emergency in Santa Barbara County on 20 May. Oiled birds and mammals have been recovered, some still alive, as government and

private workers try to determine the cause of the leak and contain the damage. According to Plains All American Pipeline, the company in Houston, Texas, that operates the pipeline, a maximum of 2,400 barrels have leaked. Around 500 barrels' worth had been recovered by the end of 24 May.

include testing the effects of the harsh space environment on various materials. The launch also sent ten miniature satellites into orbit, as well as a privately funded solar sail, 'LightSail'.

Surgeon misconduct

Paolo Macchiarini, famed for transplanting artificial windpipes into patients, committed scientific misconduct in six papers concerning the procedure, according to an independent investigator. Bengt Gerdin at Uppsala University in Sweden found that Macchiarini's papers made the operation sound more successful than it was. The Karolinska Institute in Stockholm, where Macchiarini is a visiting professor, asked Gerdin to examine allegations of scientific misconduct

against Macchiarini after four physicians who had cared for three of the trachea recipients filed complaints. Macchiarini now has two weeks to formally respond to the report. See page 406 for more.

RESEARCH

Gibbon rescue plan

A plan to save the world's rarest primate, which resides on China's Hainan Island, has been produced by an international team led by the Zoological Society of London. The plan, released on 19 May, cites a need to limit human activity in the 20 square kilometres of forest that contain the last 25 Hainan gibbons (*Nomascus hainanus*), as well as increased monitoring and habitat improvements. If it becomes extinct, the gibbon

would be the first ape wiped out by human activity: much of its decline from numbers of around 2,000 in the 1950s has been driven by habitat loss from logging (see *Nature* **508**, 163; 2014).

Dolphin deaths

The deaths of more than 1,300 bottlenose dolphins beached in the northern Gulf of Mexico since early 2010 may be related to the massive Deepwater Horizon oil spill, according to a study published on 20 May (S. Venn-Watson *et al.* *PLoS ONE* **10**, e0126538; 2015). The spike in dolphin deaths began shortly before the spill in April 2010, which has perplexed scientists. The study found that the dead animals had a much higher incidence of lung and adrenal lesions,

consistent with exposure to petroleum, compared with dolphins that were stranded before the spill or in different areas. See go.nature.com/ibmbxr for more.

Imperilled EU birds

Loss and degradation of natural habitats such as grasslands and wetlands is threatening 32% of about 450 wild bird species common in Europe. The most extensive report yet in the series *The State of Nature in the European Union*, by the European Environment Agency, finds that species including once-common farmland birds such as the skylark (*Alauda arvensis*) are at risk of extinction or in critical decline. Over-grazing and widespread use of fertilizers and pesticides are among the main threats to bird habitats. However, thanks to targeted conservation, about half of wild bird species are deemed secure.

Amazon railroad

Brazil and China agreed on 19 May to launch feasibility studies on a transcontinental railroad that would cut transportation costs for soya beans and other exports. China, Brazil and Peru signed a memorandum of agreement last year on the 5,300-kilometre railroad, which would cut through the Amazon and connect ports on the Atlantic and Pacific coasts.



Chinese Premier Li Keqiang (pictured, centre) made the 'Twin Ocean Railroad', estimated to cost around US\$10 billion, a cornerstone of his visit to Latin America. Activists are already raising concerns about impacts on the rainforest as well as on indigenous tribes.

Retraction request

The senior author on a much-publicized *Science* paper on changing people's views about same-sex marriage asked for it to be retracted on 19 May. Donald Green at Columbia University, New York, made the request after other researchers spotted irregularities in data provided by his co-author Michael LaCour, a graduate student at the University of California, Los Angeles. The study found that short conversations with a canvasser who is gay could encourage people to support same-sex

marriage. *Science* posted an 'expression of concern' online on 20 May, noting that the request was made "because of the unavailability of raw data and other irregularities".

PEOPLE

Espionage arrests

Two professors are among six Chinese men who were charged with economic espionage by the US Department of Justice on 19 May. They are accused of stealing information from US companies about an acoustic chip that has consumer electronics and military applications. On 16 May, one of the accused, Zhang Hao, an engineering professor at Tianjin University in China, was arrested in Los Angeles, California, as he got off a flight to attend a conference. This follows other recent cases in which the United

COMING UP

30 MAY – 2 JUNE

The latest microbial research is presented to about 8,000 researchers from around the world at the meeting of the American Society for Microbiology in New Orleans, Louisiana. Topics include biomedicine, food safety and environmental microbiology. go.nature.com/bqdgs4

1–5 JUNE

Physicists from a range of fields meet at the Kavli Institute for Theoretical Physics at the University of California, Santa Barbara, for the Closing the Entanglement Gap conference. It focuses on how entanglement can elucidate problems in quantum information, holography and condensed matter systems. go.nature.com/nrk7xe

States has levelled accusations of economic espionage at scientists for the benefit of China.

John Nash dies

US mathematician John Nash and his wife Alicia were killed in New Jersey on 23 May when their taxi was involved in an accident. Nash won the 1994 Nobel economics prize for a breakthrough he had made in 1950 on game theory. In March, he was awarded an Abel Prize for his contributions to the theory of nonlinear partial differential equations and its applications to geometric analysis. For much of his life, Nash, who was 86, struggled with schizophrenia, an experience portrayed in the 2001, Oscar-winning film *A Beautiful Mind*.

NATURE.COM

For daily news updates see: www.nature.com/news

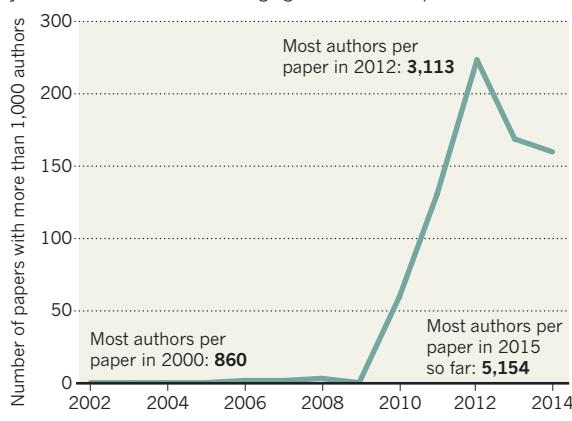
TREND WATCH

SOURCE: THOMSON REUTERS SCIENCEWATCH

More papers with huge lists of authors have appeared in the past few years, according to data from Thomson Reuters. Authorships of more than 1,000 were very rare in the early 2000s, but by 2011 more than 130 such papers had been published. The size of these multi-contributor lists has also increased, reaching the record-breaking figure of 5,154 this month for a paper in *Physical Review Letters* by two teams at Europe's particle physics lab, CERN (see *Nature* <http://doi.org/4sn>; 2015).

MEGA-AUTHORSHIPS SOAR

The number of papers with vast contributor lists has climbed in recent years — as has the size of these gargantuan authorships.



NEWS IN FOCUS

MICROBIOLOGY Array of antibiotic alternatives could fight resistant microbes **p.402**

CLIMATE CHANGE Carbon reduction pledges fall short of UN goals **p.404**

ARTIFICIAL WINDPIPES Famous surgeon committed scientific misconduct **p.406**



DEFENCE Laser weapons go from science fiction to battlefield fact **p.408**

MICHAELA REHLE/REUTERS/CORBIS



Activists campaign against a trade treaty that they fear will allow US chlorine-washed chicken into the EU.

INTERNATIONAL LAW

Trade treaties centre on science

Proposed deals have potential to boost research, but also to weaken health and environmental protections.

BY DANIEL CRESSEY

Two treaties that would govern most of the world's trade — and change how nations across the globe use scientific evidence to craft regulations — inched closer to fruition this week. On 22 May, the US Senate approved legislation that could speed up approval of the Transatlantic Trade and Investment Partnership (TTIP) and the Trans-Pacific Partnership (TPP).

Public attention has focused on the

economic impact of the treaties: TTIP is between the United States and the European Union (EU) and the TPP is between a variety of North American and Asian nations (including the United States, Mexico, Japan, Australia and Malaysia, but not China or India) whose economies together account for around 60% of the world's gross domestic product. But researchers, industry groups and non-governmental organizations (NGOs) are urging scrutiny of the treaties' impact on science.

Supporters point to the agreements'

potential to boost research in pharmaceuticals. Critics say that the treaties could undermine countries' abilities to protect the health of their citizens, as well as to cleave to nation-specific comfort levels with genetically modified (GM) organisms.

"TTIP really falls squarely within the domain of science and technology policy," says Sebastian Pfotenhauer, a science-policy researcher at the Massachusetts Institute of Technology in Cambridge. "What is at stake is the sovereignty of countries to interpret scientific data and regulate risks in the way they choose."

In the case of TTIP, tariffs between the EU and United States are already so low that the benefits to business are expected to come mainly from harmonization of regulations, which differ significantly between the two blocs. The EU is sceptical of GM crops, limiting what can be grown and imported there. In the United States, by contrast, the crops are widespread and are generally accepted by the public. However, the United States has tougher standards in other areas of trade: for example, it heavily restricts beef purchases from the EU because of fears over BSE (bovine spongiform encephalitis), also known as mad cow disease.

RACE TO THE BOTTOM

Some groups fear that agreements such as TTIP could be used to force all nations to adopt the rules of the country with the least regulation in each field. "It sets up a series of regulatory steps that will undoubtedly create a race to the bottom," says David Azoulay, an attorney with the Center for International Environmental Law (CIEL) in Geneva, Switzerland. "We believe that there are basically no benefits that will come from TTIP for people or the environment," adds Fabian Flues, a campaigner for Friends of the Earth Europe in Geneva. CIEL notes that the agrochemical industry is already considering using TTIP to increase the acceptable level of pesticide residues on food by pushing for the higher US threshold to be adopted in the EU.

Others dismiss the idea of a 'race to the bottom'. The European Commission says that harmonization need not lower regulatory protection. And Alberto Alemanno, a lawyer at the business school HEC Paris who has done consultancy work for the EU's TTIP negotiating team, says that TTIP aims to create a system that allows regulators to discuss when ▶

► their standards are equivalent, which does not necessarily mean making all regulations identical, as many people assume. If regulations are harmonized in some instances, he points out, citizens would be able to push for the higher standard — not the lower — to be adopted.

Greater clarity on how medicines will be treated by regulators would be a boon for pharmaceutical research, says Richard Bergström, director-general of the European Federation of Pharmaceutical Industries and Associations in Brussels. It could also mean that drugs reach patients earlier, because companies would no longer have to apply separately to US and EU regulators, a process that can involve costly and time-consuming

tests, especially for paediatric medicines. "I'm both surprised and quite frustrated by the debate around TTIP," he says.

Another fear surrounding the trade agreements is that if they are ratified they

will include a legal framework known as an investor-to-state dispute resolution mechanism, which would allow companies to take governments to court and overturn legislation. A company could thus limit a government's ability to regulate a certain chemical on health grounds if that were seen as an unreasonable restriction of the company's trade as laid out in the treaty. Objectors to the treaties fear that a dispute provision would also discourage governments from passing strict legislation in the first place.

The history of the North American Free Trade Agreement (NAFTA) and other such agreements is littered with such cases, says Pfotenhauer. The pharmaceutical firm Eli Lilly is currently suing Canada under NAFTA after the government invalidated patents on two of the company's drugs in an argument over interpretation of clinical-trial data. Companies have also tried to use NAFTA to force Canada to allow toxic-waste exports and to include an additive in petrol that the nation's regulators claimed was dangerous.

These cases revolved around the interpretation of scientific knowledge, notes Pfotenhauer. Indeed, one thing that NGOs, lawyers and governments agree on is that scientists should care about the trade deals. "TTIP is set to change the way in which scientific cooperation is set to take place across the Atlantic," says Alemanno.

And there may be little point in fighting the deals. Against a background of gloom and pro-free-trade governments, says Alemanno, even if this particular trade negotiation breaks down, as a philosophy, "TTIP is here to stay". ■ SEE EDITORIAL P.393

"TTIP is set to change the way in which scientific cooperation is set to take place."

DRUG DEVELOPMENT

Bacterial arms race revs up

With antibiotic resistance on the rise, researchers are looking for new ways to treat infections.

BY SARA REARDON

More than eight decades have passed since Alexander Fleming's discovery of a fungus that produced penicillin — a breakthrough that ultimately spawned today's multibillion-dollar antibiotics industry. Researchers are now looking to nature with renewed vigour for other ways of fighting infection.

Few new antibiotics are in development, and overuse of existing ones has created resistant strains of deadly bacteria. "We need a change from what we have," says Stephen Baker, head of medicinal chemistry for antibacterials at GlaxoSmithKline in Collegeville, Pennsylvania.

Baker will talk about some of the alternatives to antibiotics on 2 June at the American Society for Microbiology's annual meeting in New Orleans, Louisiana. Here are a few of the therapies that scientists are exploring.

PREDATORY BACTERIA

Bacteria cause infection, but some can also fight it by preying on fellow microbes. Several researchers are beginning to test these predatory bacteria in animal models and cell cultures.

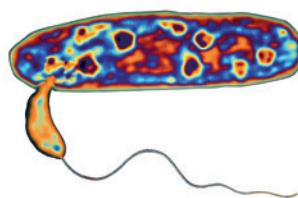
The best-known species, *Bdellovibrio bacteriovorus*, is found in soil. It attacks prey bacteria by embedding itself between the host's inner and outer cell membranes, and begins to grow filaments and replicate. "It's like going

into a restaurant, locking the door and starting to munch away," says Daniel Kadouri, a bacteriologist at Rutgers University in Newark, New Jersey. The host bacterium eventually explodes and releases more *B. bacteriovorus* into the environment.

Kadouri and others are also studying the therapeutic potential of the predatory bacterium *Micavibrio aeruginosavorus*. And a team has engineered the gut bacterium *Escherichia coli* to produce peptides that kill *Pseudomonas aeruginosa*, a microbe that causes pneumonia.

This preliminary research is attracting attention. The Pathogen Predators programme of the US Defense Advanced Research Projects Agency, which aims to treat soldiers who contract infections on the battlefield, announced nearly US\$16 million in research grants this week to groups studying predatory bacteria.

ALFRED PASIEKA/GETTY



Bdellovibrio bacteriovorus preys on other bacteria.

ANTIMICROBIAL PEPTIDES

Plants, animals and fungi have vastly different immune systems, but all make peptides — small proteins — that destroy bacteria. Peptides from creatures such as amphibians and reptiles, which are unusually resistant to infection, could yield new therapeutics.

Peptides with antibacterial activity have been isolated from frogs, alligators and cobras, among others, and some seem to be effective in epithelial cell cultures and at healing wounds in mice. These peptides can be modified to increase their potency, and several are

**MORE
ONLINE**



TOP STORY



Global ocean trawl reveals plethora of new life-forms go.nature.com/2sphxb

MORE NEWS

- Alzheimer's origins tied to rise of human intelligence go.nature.com/3ckf7k
- Ancient wolf genome pushes back dawn of the dog go.nature.com/zltqia
- 'Stable' region of Antarctica is melting go.nature.com/gt0hrs



Peptides produced by the American alligator could add to the arsenal of antibacterials.

in clinical trials. One, called pexiganan, based on a peptide from frog skin, is now in phase III clinical trials to treat diabetic foot ulcers.

But synthesizing such molecules can be expensive, a hurdle that scientists must overcome to bring new peptide drugs to market.

PHAGES

Of all the alternatives to antibiotics, phages — viruses that attack bacteria — have been used the longest in the clinic. Scientists in the Soviet Union began developing phage therapies in the 1920s, and former Soviet countries continue the tradition.

Phages have several advantages over antibiotics. Each type attacks only one type of bacterium, so treatments leave harmless (or beneficial) bacteria unscathed. And because phages are abundant in nature, researchers have ready replacements for any therapeutic strain that bacteria evolve to resist.

Mzia Kutateladze, who heads the scientific council at the Eliava Institute in Tblisi, Georgia, says that antibiotic resistance is driving more Western patients to phage-therapy clinics in Eastern Europe. The US National Institute of Allergy and Infectious Diseases in Bethesda, Maryland, now lists phages as a research priority for addressing the antibiotic crisis. A clinical trial of a phage treatment for infections associated with burns is planned by a consortium of European centres to start this summer.

GENE-EDITING ENZYMES

CRISPR, a gene-editing technique that has taken the scientific world by storm, is based on a strategy that many bacteria use to protect themselves against phages. Researchers are turning that system back on itself

to make bacteria kill themselves.

Normally, the bacteria detect and destroy invaders such as phages by generating a short RNA sequence that matches a specific genetic sequence in the foreign body. This RNA snippet guides an enzyme called Cas9 to kill the invader by cutting its DNA.

Scientists are now designing CRISPR sequences that target genomes of specific bacteria, and some are aiming their CRISPR kill switches at the bacterial genes that confer antibiotic resistance.

METALS

Metals such as copper and silver are the oldest antimicrobials. They were favoured by Hippocrates in the fourth century BC as a treatment for wounds, and were used even earlier by ancient Persian kings to disinfect food and water. Only now are researchers beginning to understand how metals kill bacteria.

Some groups are exploring the use of metal nanoparticles as antimicrobial treatments, although little research has been done in people. Because metals accumulate in the body and can be highly toxic, their use may be restricted mostly to topical ointments for skin infections.

An exception is gallium, which is toxic to bacteria that mistake it for iron, but is safe enough in people to be tested as an intravenous treatment for lung infections. This summer, researchers at the University of Washington in Seattle will begin a phase II clinical trial of gallium in 120 patients with cystic fibrosis. Pilot studies found that the metal was moderately successful at breaking down microbial biofilms in the lungs and improving patients' breathing. ■



China accounts for almost one-third of global carbon dioxide emissions.

CLIMATE CHANGE

Scientists step in to assess carbon-emissions pledges

Lack of common standard for national climate plans puts focus on independent ratings.

BY JEFF TOLLEFSON

The road to Paris is paved with promises. In the run-up to the United Nations climate summit in the French capital this December, nations are submitting pledges to limit their future carbon emissions. With a preliminary session in Bonn, Germany, scheduled for 1–11 June, it is already clear that the commitments are insufficient to meet the UN's goal of limiting future temperature increases to 2 °C above pre-industrial levels. It is difficult to say which countries should do more, but in an effort to move the process forward, climate scientists have started to independently assess the pledges. They have found that a Paris agreement to limit warming to 2 °C would require unrealistic political compromises.

The big picture is clear: on the basis of pledges made so far, the global average temperature is projected to rise by around 3 °C above pre-industrial levels by 2100. For a reasonable chance of meeting the 2 °C goal, countries would have to strengthen their pledges sufficiently to limit all future greenhouse-gas emissions to the equivalent of about 1 trillion tonnes of carbon dioxide.

"Once you say we are not doing enough, it begs the question, 'Who should do more?'" says Joe Aldy, an economist at Harvard

University in Cambridge, Massachusetts, and former special assistant for energy and environment to US President Barack Obama.

That is the issue that climate scientists are trying to grapple with — without taking sides. "This whole issue of assessing these national targets is an important one, and it's definitely a place where science can help," says Niklas Höhne, a founding partner of the NewClimate Institute in Cologne, Germany. "But equity and fairness is something which is very much up to interpretation — what's fair for one is not fair for another."

The NewClimate Institute is part of a research consortium that produces the Climate

"It's going to be a little bit messy this year."

Action Tracker, a comprehensive assessment of the pledges made thus far. Using the peer-reviewed literature as a source, it rates pledges on the basis of how close they would bring the world to the 2 °C goal if all countries adopted similar standards.

The United States, European Union (EU) and China get a 'medium' rating from the Climate Action Tracker, meaning that their level of commitment, if adopted globally, would probably lead to more than 2 °C of warming. In a pledge submitted on 31 March, the United States vowed

to reduce its emissions to at least 26% below 2005 levels by 2025. China has not submitted an official pledge, but committed in November to start reducing its emissions around 2030, the same year that the EU plans to have cut emissions to 40% below 1990 levels.

Canada, which pledged on 15 May to reduce its emissions to 30% below 2005 levels by 2030, received an 'inadequate' rating from the tracker, meaning that a global application of its plans would lead to warming of 3–4 °C or more.

REPORT CARDS

Criteria are key for assessing these pledges. In a hypothetical agreement, warming could be limited to 2 °C by allocating shares of the remaining emissions budget among nations. But different methods can be used to allot emissions: nations' proportions would differ depending on whether shares were distributed on the basis of emissions per capita, per dollar of gross domestic product, or according to the current distributions (see 'Cutting the carbon cake'). In general, developing countries contend that those responsible for a greater share of historical emissions should get a smaller share of the future emissions budget. Rich countries counter that most emissions today and in the future will come from developing countries, which must also do their part.

SOUVID DATTA

A scientific framework published in *Nature Climate Change* last September created two emissions-quota scenarios based on the contrasting positions: one favoured poorer countries by allotting emissions on a per capita basis; the other preserved the status quo by carrying forward countries' current shares of emissions (M. R. Raupach *et al.* *Nature Clim. Change* **4**, 873–879 (2014)). The study also produced a third, 'compromise' approach. The authors found that this scenario worked best to meet the 2 °C target — but it would require politically and economically difficult shifts in energy use by all countries.

A co-author of that study, Glen Peters of the Center for International Climate and Environmental Research in Oslo, has since analysed the US, EU and Chinese commitments using the same framework. His unpublished analysis finds that the US and EU pledges would eat up a large proportion of future emissions to meet the 2 °C goal, leaving developing countries with a relatively small share going forward.

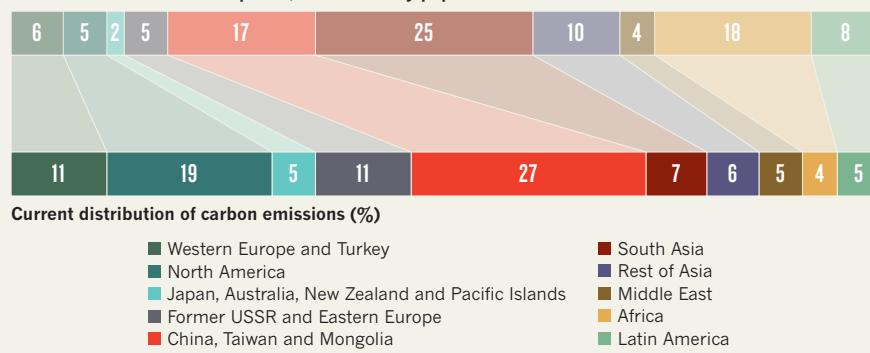
Poorer countries would surely balk at such a plan; they want room to grow and they see cheap fossil-fuel energy as a legitimate tool to lift their citizens out of poverty. Although it has yet to submit its pledge, India, for instance, has said that it will not limit its CO₂ emissions until it has risen to the global average of around 5.1 tonnes per capita, from its current level of around 1.9 tonnes per capita.

And then there is China. It now accounts for around 27% of global CO₂ emissions and its per capita emissions are 6% higher than the EU's. Peters' calculations show that the emissions reductions required of China would be

CUTTING THE CARBON CAKE

To keep global warming to within 2 °C above pre-industrial levels, the world must keep future carbon emissions to about 1 trillion tonnes. Allocating shares of that total to nations on the basis of population size results in a dramatically different allocation of emissions compared with carrying forward current distributions.

Potential future emissions quotas, if allocated by population size



so drastic as to be politically unrealistic: even under the most lenient terms it would have to halt its rapidly growing emissions this year and cut them almost to zero by 2050. In that light, Peters says that climate scientists need to focus more on political viability. "The scientific community in a sense might be a little bit unprepared for doing this kind of thing."

China and India issued a joint statement on 15 May that urged developed countries to make bigger emissions cuts before 2020, when a Paris agreement would probably take effect. The statement also asked developed countries to honour their pledge to provide US\$100 billion per year in aid by 2020 to help poor countries to adapt to warming and reduce emissions.

Ultimately, Aldy says, the UN climate process needs to incorporate a formal review similar to those used by other international bodies, such as the World Trade Organization and the International Monetary Fund. At those organizations, experts produce peer-reviewed reports on economic and trade policy using standardized reports of data and regulations submitted by member countries. The scientific analyses being done now may represent a small step in that direction.

"It's going to be a little bit messy this year," Aldy says. "The question is whether that means we end up having a bunch of unproductive conversations about what countries are doing, or whether this turns out to be a productive process of experimentation." ■

PUBLIC HEALTH

Ebola R&D woes spur action

Public-health officials make plans for how to speed up research-and-development responses to tropical-disease outbreaks.

BY ERIKA CHECK HAYDEN

After failing to deliver new treatments in time to fight the now-waning Ebola epidemic in West Africa, public-health experts are planning for speedier clinical trials during future outbreaks.

"Everyone is delighted that we are going to no cases," says the World Health Organization's (WHO) assistant director-general Marie-Paule Kieny, "but research and development is too late by a factor of months. We need to see what can be done in order to have the products more rapidly available for another outbreak."

In an interview last week at the annual

WHO meeting of member states in Geneva, Switzerland, Kieny said that the agency is working on a framework to identify pathogens that merit advanced research. The framework would catalogue the work needed to prepare potential treatments, vaccines and diagnostics for large-scale efficacy testing in the event of an outbreak.

The plan, a draft of which could be ready by the end of the year, will also spell out how researchers, companies and health agencies should agree to share data collected in publicly funded studies.

It was only in March — a full year after the Ebola outbreak was first reported — that large-scale tests of candidate vaccines got under way.

Some of those trials have little chance of being completed because cases are now too rare: the US National Institute of Allergy and Infectious Diseases, which is testing a preventive vaccine, estimates that it would need to enrol more than 150,000 people in Guinea to show that its vaccine works.

Although trials of Ebola drugs and vaccines were set up at a breakneck pace compared with the years it normally takes to organize large clinical trials, the researchers behind them have been frustrated by delays. "We have said that we need to conduct trials in this outbreak, and we've largely failed, and that's desperately disappointing," says Trudie Lang, a global-health

► researcher at the University of Oxford, UK, who has been involved with clinical trials of Ebola treatments.

Similar concerns arose during the severe acute respiratory syndrome outbreak of 2003–04 and the 2009 H1N1 flu pandemic; after both events, researchers drafted study designs that could speed up future trials in outbreaks caused by respiratory pathogens. But there has not been a similar effort to coordinate trials for haemorrhagic diseases such as Ebola, which are less likely to spread to wealthy nations.

The WHO framework aims to boost readiness to conduct trials for diseases that could cause global health emergencies. Although it has not been decided which diseases will qualify, Lassa fever, Rift Valley fever and Middle East respiratory syndrome are candidates.

DIFFERENT PRIORITIES

The delays in developing trials of Ebola treatments, researchers say, were due to a disconnect between organizations caring for patients and those leading the trials; refusals to share data; and failures to perform the efficacy and safety tests that could have been done before the outbreak (see *Nature* 511, 520; 2014).

“We need to do the phase I studies in humans and have a small stockpile rather than stop at the animal stage,” says England’s chief medical officer Sally Davies.

Trials also lagged because of disagreements over logistics and ethics. For example, after the medical aid group Médecins Sans Frontières (Doctors without Borders) and researchers at the University of Oxford organized trials of the drug brincidofovir that did not include placebo groups, the US Food and Drug Administration pushed for a randomized trial that would include untreated controls. The trial eventually went forward without a control group.

Data-sharing roadblocks have also occurred at all levels. The experimental drug ZMapp was given to a handful of patients before supplies ran out in August 2014. But detailed information on the patients’ reactions to the drug has not been released, owing to fears that this would prevent researchers from publishing on the cases. Similarly, pharmaceuticals firm Chimerix of Durham, North Carolina, would not publicly reveal why it withdrew support for a trial of brincidofovir in late January after four patients had been treated, but Kiely says that the decision was due to the drug’s lack of efficacy.

At a broader scale, the WHO cannot publicly release comprehensive epidemiological information because the data are owned by countries where cases occur.

Kiely is trying to persuade publishers to agree that researchers can release analyses of patient and laboratory data during health emergencies without compromising their ability to publish on it. She is also trying to broker an agreement with funding agencies that data from publicly funded safety and efficacy trials in outbreaks must be made widely available. ■



YURI BELINSKY/TASS/CORBIS

Paolo Macchiarini has misrepresented data on the success of artificial-trachea transplants, says a report.

SYNTHETIC WINDPIPES

Surgeon commits misconduct

Papers authored by Paolo Macchiarini misrepresented success of pioneering tracheal transplant procedure.

BY DAVID CYRANOSKI

A surgeon famed for his pioneering transplants of synthetic windpipes has committed scientific misconduct, according to an independent investigation.

Six published papers authored by thoracic surgeon Paolo Macchiarini, a visiting professor at the Karolinska Institute in Stockholm, had misrepresented data from recipients of the artificial windpipes, or tracheas, reports Bengt Gerdin, a general surgeon and professor emeritus at Uppsala University in Sweden. The papers made the operation sound more successful than it was, says Gerdin, who was commissioned by the prestigious Karolinska Institute to examine Macchiarini’s clinical procedures.

Gerdin also found that two of the papers^{1,2} described operations that had not received the necessary ethical approval, and that a seventh paper³ authored by Macchiarini, reporting transplants of artificial oesophagi into rats, had misrepresented results.

Macchiarini told *Nature* that he would

not comment on the investigation report, which is in Swedish, until he had seen an English version. In November 2014, he told *Nature* that he welcomed the investigation and was confident that “there is nothing suspect, unethical, inflated or misleading about anything I have done or reported”. He has two weeks to formally respond to Gerdin’s report.

His research had been hailed as a bright spot for regenerative medicine, which has been slow to deliver synthetic materials to replace natural organs. Macchiarini’s procedure involved bathing a polymer trachea in stem cells from the transplant recipient’s bone marrow. The idea was that when this was used to replace a damaged trachea, the cells would form the right type of tissue and create a seal with the surroundings.

Macchiarini has put such artificial tracheas into eight people. The papers under investigation relate to just three of the procedures, and report that there were some signs that the synthetic tracheas had successfully integrated. Two of those transplant recipients

have died; one entered intensive care after the procedure and is still there. Macchiarini has previously told *Nature* that the problems faced by the patients were unrelated to the transplants.

The investigation began after four Karolinska physicians who were involved in the care of those three patients filed complaints. Karl-Henrik Grinnemo, Matthias Corbascio, Thomas Fux and Oscar Simonson provided medical records that they alleged to be at odds with the published results, and called into question the paper on the rat model.

In compiling the report, Gerdin says, he tried to avoid matters of interpretation that would become “a quarrel between scientists”, and stuck to facts such as whether the medical records showed evidence of a follow-up at the intervals claimed. In some instances, publications claimed improvement even though there was no evidence that the patients had been examined. “This is falsification,” says Gerdin. Speaking of Macchiarini, he adds: “The basic rule in science is to have all reports documented, but he doesn’t have them.” The rat-model paper included weight-gain and computed-tomography (CT) data that had been misinterpreted to suggest that the oesophageal graft was more successful than it was, says Gerdin.

He concludes that the misrepresentations were deliberate: “If there is a mistake once, you might think it is random. If it happens several times, you begin to question whether it really is random.”

The investigation focused on Macchiarini, but Gerdin notes that Grinnemo, Corbascio and Simonson were co-authors of one of the papers² included in the investigation involving a patient, and that Grinnemo was also a co-author of a second¹. “I’m not saying they share the responsibility, but one has to ask what they knew and what

they didn’t know,” says Gerdin.

The three physicians did not respond to a request for comment, but Corbascio had previously told *Nature* that he was involved only “at a superficial level” with the transplant recipient. “I had complete confidence in Macchiarini,” he said.

Separately, on 9 April, the Swedish Medical Products Agency (MPA) filed a complaint with the Swedish state prosecutor over whether proper permission was obtained to carry out the three synthetic tracheal transplants that feature in Gerdin’s investigation. The operations took place at the Karolinska. Ann Marie

Janson Lang, a clinical assessor in the clinical-trials department of the MPA, says that the synthetic tracheas meet the definition of an “advanced therapy medicinal product”, which requires agency permission before it can be given to patients, but that no application for a permit was made.

It is not clear who would bear responsibility for the breach, if it is confirmed. Macchiarini told *Nature* that, as a visiting professor at the Karolinska, “it was never my responsibility to obtain any necessary permissions. I was not directed to do so, nor did I have the authority to do so.” As *Nature* went to press, a Karolinska spokesperson said that the institute was preparing to post on its website “the facts about the three operations that Macchiarini performed and other matters connected to that case”. The state prosecutor is expected to report back within weeks. ■

1. *Lancet* **378**, 1997–2004 (2011).

2. *Biomaterials* **34**, 4057–4067 (2013).

3. *Nature Commun.* **5**, 3562 (2014).

CORRECTION

The map in the News Feature ‘India by the numbers’ (*Nature* **521**, 142–143; 2015) omitted the province of Meghalaya. The full map can be seen at go.nature.com/h2ydqb.



LASER WEAPONS GET REAL

Long a staple of science fiction, laser weapons are edging closer to the battlefield — thanks to optical fibres.

BY ANDY EXTANCE

Silently, the drone aircraft glides above the arid terrain of New Mexico — until it suddenly pivots out of control and plummets to the ground.

Then a mortar round rises from its launcher, arcs high and begins to descend towards its target — only to flare and explode in mid-flight.

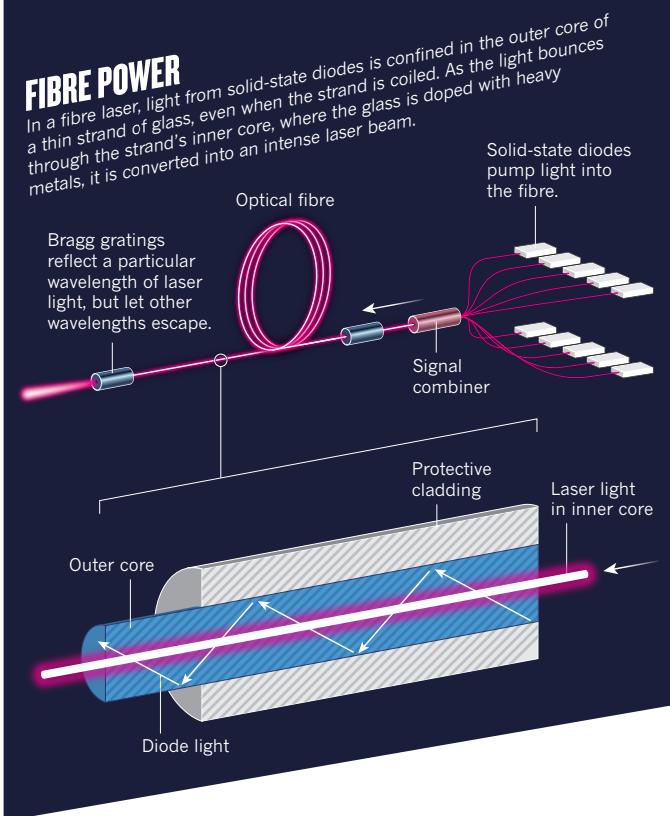
On the desert floor, on top of a big, sand-coloured truck, a cubic mechanism pivots and fires an invisible infrared beam to zap one target after another. This High Energy Laser Mobile Demonstrator (HEL MD) is a prototype laser weapon developed for the US Army by aerospace giant Boeing of Chicago, Illinois. Inside the truck, Boeing electrophysics engineer Stephanie Blount stares at the targets on her laptop's screen and directs the laser using a handheld game controller. "It has a very game-like feel," she says.

That seems only natural: laser weapons are a staple of modern video

games, and ray-guns of various sorts were common in science fiction for decades before the first real-life laser was demonstrated in 1960. But they are not a fantasy anymore. The Boeing prototype is just one of several such weapons developed in recent years in both the United States and Europe, largely thanks to the advent of relatively cheap, portable and robust lasers that generate their beams using optical fibres.

The output of these fibre weapons is measured in kilowatts (kW), orders of magnitude less than the megawatt-class devices once envisioned for the US Strategic Defense Initiative — an ultimately unsuccessful cold-war plan that sought to use lasers to disable ballistic missiles carrying nuclear warheads.

But the modern, less ambitious, weapons are on the brink of real-world deployment. Tests such as those of the Boeing system show that the lasers have enough power to overcome threats from terror



groups — at a fraction of the price of conventional defences. “It’s a very cost-effective solution to taking out cheaply made weapons like small mortars or rockets made out of sewer pipe,” says Blount.

In late 2014, for example, the US Navy showed that a ship-mounted laser-weapon system called LaWS could target small boats, such as those used by terrorists and pirates. That experimental weapon is currently installed on the USS *Ponce*, an amphibious support ship in the Gulf.

Many challenges to full-scale deployment remain, warn developers, from the need to boost the weapons’ power to the difficulty of operating a laser in fog and clouds. But specialists in defence and security are starting to take lasers seriously. “After a nearly half-century quest, the US military today is on the cusp of finally fielding operationally relevant directed-energy weapons,” wrote Paul Scharre, an advanced-technology specialist at the Washington DC-based Center for a New American Security (CNAS), in report on laser weapons released in April¹.

THE POWER PREDICAMENT

Laser weapons have long fascinated weapons developers — most notably during the heyday of the Strategic Defense Initiative, nicknamed Star Wars, in the 1980s and 1990s. US spending on laser-weapons research peaked in 1989 when, according to the CNAS report¹, the government spent the equivalent of US\$2.4 billion in 2014 dollars. Funding has continued at lower levels ever since. Yet the original goal, of being able to shoot down incoming ballistic missiles, proved unattainable.

The trick with any laser weapon is to focus its energy into a spot that is small enough to heat up and damage the target — and to do that with a machine that is compact and portable enough for the battlefield. This is easier said than done. In 1996, for example, the US Air Force initiated the Airborne Laser project as one of its contributions to defence against ballistic missiles. Because it was impossible at the time to generate the required megawatts of optical power electrically, the developers chose a chemical oxygen iodine laser (COIL) that could be fuelled by a chemical reaction. But the COIL was so bulky that it could only be carried on a Boeing 747, and left little space for laser fuel. “It needed remote mixing units and chemicals weighing tens of thousands of pounds,” says Paul Shattuck, head

NATURE.COM

For a podcast on
laser weapons, visit:
go.nature.com/2xdtw4

of directed-energy systems for Lockheed Martin Space Systems, which provided the project’s beam-control technology.

Another major problem was the atmosphere, says Phillip Sprangle, senior scientist for directed-energy physics at the Naval Research Laboratory in Washington DC. Not only was the beam scattered by dust and natural turbulence, he says, but its passage caused ‘thermal blooming’. When the beam propagated at very high powers, Sprangle explains, “the atmosphere absorbed laser light, heating the air and causing the laser beam to spread out”. That spreading, in turn, dissipated the laser’s energy.

The good news for the Airborne Laser project was that this issue, at least, had a solution: adaptive optics technology similar to that used by astronomers to clarify their view of the stars (see *Nature* 517, 430–432; 2015). The technology uses mirrors to automatically distort the laser beam in a way that cancels the effects of the turbulence, with the same result as a pair of glasses correcting for aberrations in the eye. “As the laser beam passes through the atmosphere,” says Shattuck, “it cleans up, and it’s nice and tight when it gets to the target.”

By 2010, the adaptive optics was good enough for the Airborne Laser to destroy a ballistic missile in flight. By then, however, logistical issues such as the size problem had led the Department of Defense to lose its enthusiasm for energy weapons in general. It cancelled the Airborne Laser programme outright by early 2012. At the same time, the department’s spending on high-energy lasers in general was falling; it dropped from \$961 million in 2007 to \$344 million in 2014.

FIBRES IN THE SPOTLIGHT

The money did not vanish entirely: attention was already shifting to fibre lasers as a way to deliver results more economically. Fibre lasers were invented in 1963, and since the 1990s they have been advanced almost entirely by IPG Photonics in Oxford, Massachusetts. Whereas other solid-state lasers use rigid rods, slabs or discs of crystal to generate the beam, and so have to be fairly large, fibre lasers use thin optical fibres that can be wrapped into compact coils (see ‘Fibre power’). The fibres can collect their optical energy from brighter versions of the cheap laser diodes used in DVD players, and then amplify the light to higher power, with overall electrical-to-optical conversion efficiencies greater than 30%. This is at least double the efficiency typical of other solid-state lasers, and close to that of chemical lasers such as COIL. And, being intrinsically long and thin, the fibres have a high surface area to volume ratio and can radiate away waste heat very quickly — an ability that helps to give the lasers a long working life and low maintenance requirements.

These advantages first attracted attention during the 1990s, when fibre lasers began to be used to beef up optical signals carrying Internet data through undersea cables. But since the early 2000s, IPG has focused on developing kilowatt-class industrial lasers for welding, drilling and cutting — devices that also attracted the attention of military researchers.

Around 2010, recalls Shattuck, he and his colleagues at Lockheed Martin heard from Israeli civilians targeted by rockets launched from the Gaza Strip. “The mayor of a village stood up and said, ‘Please, give me some kind of defence,’” Shattuck says. This inspired Lockheed Martin to develop the Area Defense Anti-Munitions (ADAM) system, which uses an off-the-shelf 10-kW laser from IPG to keep costs down. Since 2012, the company has shown that ADAM can disable targets such as boats, drones and simulated small-calibre rockets from about 1.5 kilometres away. Although unwilling to disclose the price of ADAM — or whether anybody has bought one — Lockheed Martin says that it is now ready to provide the system to customers.

Blount is less reticent about Boeing’s HEL MD prototype, which also uses a commercial 10-kW fibre laser. With the system drawing its power from the vehicle engine or a separate generator, she says, “it takes less than two cups of fuel to fire the laser for long enough to disable many targets.” This makes it much cheaper to use for defence than conventional missiles. “An inexpensive missile is \$100,000 and that’s one shot,” says David DeYoung, Boeing’s director of directed-energy systems. “To shoot a laser-weapon system once is less than \$10.”

Blount stresses that the resurgence of laser weapons owes at least as

much to advanced image-recognition and targeting systems as to the laser itself. "The better the pointing and tracking system," she says, "the better able you are to put the beam on the most vulnerable point of a target."

Thanks to computerized aiming, HEL MD can operate in wholly autonomous mode, which Boeing tested successfully in May 2014 — although the trials uncovered an unexpected challenge. The weapon's laser beam is silent and invisible, and not all targets explode as they are destroyed, so an automated battle can be over before operators have noticed anything. "The engagements happen quickly, and unless you're staring at a screen 24–7 you'll never see them," Blount says. "So we've built sound in for whenever we fire the laser. We plan on taking advantage of lots of *Star Trek* and *Star Wars* sound bites."

STRENGTH IN NUMBERS

Aiming and targeting may be battle-ready, but power is still a problem. A commercial laser's 10-kW output is at the low end of what is useful for laser weapons. And using fibres puts limits on the beam's power and quality — not least because at high powers, the cascade of photons surging through the fibre can heat it up faster than it can radiate the energy, and can thus cause damage. To avoid this, researchers are working to combine the output from several lasers.

The ideal way to do this would be 'coherent combining', in which the waves from each laser march together in tightly synchronized formation. This technique is widely used in radio and microwave applications, says Tso Yee Fan, a laser scientist at the Massachusetts Institute of Technology's defence-oriented Lincoln Laboratory in Lexington. But coherence is much tougher to achieve with visible and infrared light. The waves from each laser must have almost identical wavelengths, the planes of their oscillations must precisely align, and the peaks and troughs of each wave must coincide. "In radio-frequency or microwaves, the wavelength's a few centimetres," Fan says. "In optics, the wavelength's around a micrometre, so being able to do those kinds of controls has been really difficult."

But that may not matter much, says Sprangle. In 2006, he and his team reported computer simulations suggesting that an 'incoherent combination' of several fibre-laser beams hitting a single spot would be almost as effective as a coherent combination². With either approach, he says, "when you're propagating over long ranges through atmospheric turbulence, you get approximately the same power on the target". In 2009, his group confirmed this theory by using mirrors to combine 4 fibre-laser beams into a 5-centimetre spot on a target more than 3 kilometres away³.

Building on Sprangle's work, the US Office of Naval Research has developed the 30-kW LaWS, which incoherently combines six commercial fibre-lasers. LaWS has been installed on the USS *Ponce* since September 2014, and has been tested on objects such as small boats and drones.

The missile-specialist MBDA Germany in Schröbenhausen has developed a similar approach⁴. In October 2012, the firm successfully used its 40-kW combined fibre-beam system to destroy model artillery shells towed through the air some 2 kilometres away. MBDA's tests have also helped to debunk the science-fiction idea that reflective armour would defend against laser weapons. They found that any dust on the mirrored surface would get burned in, and lead to the destruction of the target even faster than with a non-reflective surface.

Markus Martinsteller from MBDA's Future Systems Directorate argues that high-precision targeting minimizes the chances of accidentally hurting bystanders while trying to shoot down targets, especially compared with conventional explosives. "There is no risk from fragmenting ammunition and we only start the irradiation when the aim point is exactly on target," he says.

Lockheed Martin is also working on laser weapons that can take on targets that are more complex or farther away than can be tackled by its low-cost ADAM system. In March, for example, the company reported that its Advanced Test High Energy Asset (ATHENA) system could disable the running engine of a small truck mounted on a test platform. ATHENA uses a similar adaptive-optics system to the Airborne Laser, coupled with Lockheed's Accelerated Laser Demonstration Initiative (ALADIN) fibre-laser system.



The High Energy Laser Mobile Demonstrator can take down drones.

BOEING

ALADIN combines the output of several fibre lasers, each with a slightly different wavelength, into a single 30-kW beam. This 'wavelength beam combining' approach originated at the Lincoln Laboratory⁵ and is similar to methods that channel Internet traffic into fibre-optic cables. Fan notes that this method is easier than coherent combining, but gives better-quality beams than incoherent combining, so it can more easily hit smaller targets from longer distances.

Jason Ellis, a visiting fellow at the CNAS and lead author of the think tank's laser-weapons report¹, says that such developments convince him that fibre-laser weapons are coming of age — and that emerging advances could take them to hundreds of kilowatts and extend their range to hundreds of kilometres.

Despite such advances, a February 2014 poll⁶ of US national-security specialists found that just one-fifth believed that directed-energy weapon technologies would be mature within a decade.

Michael Carter, a programme manager for photon science at Lawrence Livermore National Laboratory in Livermore, California, cautions that today's lasers are a very long way from their science-fiction counterparts. "They're not yet the *Star Trek* phaser," he says. "People talk about speed-of-light engagement, but it still takes time to demolish targets. At the most basic level, if you can't see it — if there's too much rain or fog — your laser can't hit it." He suggests that the greatest value of the current generation of demonstration systems may be in working out how to handle such broader challenges before better lasers emerge. "Don't mistake what they're doing on the USS *Ponce* for a new strategic superiority," Carter warns. "It may be the first step in that direction but it's not going to change the game by itself."

Even the weapons companies are cautious not to overstate their case. For example, MBDA expects that it will take 3–5 years for truly operational systems to appear even in the tens of kilowatts range. And in some circumstances — such as a foggy day — conventional weapons will always be more effective. "You give the defender of the future both, and put the choice in their hands," DeYoung recommends.

Despite their modest capabilities, Scharre claims that fibre-laser weapons could find a niche in US military defence in 5–10 years. "They may not be as grand and strategic as the Star Wars concept," he says, "but they could save lives, protect US bases, ships and service members." ■

Andy Extance is a freelance writer in Exeter, UK.

- Ellis, J. D. *Directed-Energy Weapons: Promise and Prospects* (CNAS, 2015); available at <http://go.nature.com/eipivo>
- Sprangle, P., Peñano, J. & Hafizi, B. J. *Directed Energy* **2**, 71–95 (2006).
- Sprangle, P., Ting, A., Penano, J., Fischer, R. & Hafizi, B. *IEEE J. Quantum Elect.* **45**, 138–148 (2009).
- Mohring, B. et al. *Proc. SPIE* **8733**, 873304 (2013).
- Daneu, V. et al. *Opt. Lett.* **25**, 405–407 (2000).
- FitzGerald, B. & Sayler, K. *Creative Disruption: Technology, Strategy and the Future of the Global Defense Industry* (CNAS, 2014); available at <http://go.nature.com/7iwlaq>



Clever fish

By revealing that fish cooperate, cheat and punish, Redouan Bshary has challenged ideas about brain evolution.

BY ALISON ABBOTT

Redouan Bshary well remembers the moment he realized that fish were smarter than they are given credit for. It was 1998, and Bshary was a young behavioural ecologist with a dream project: snorkelling in Egypt's Red Sea to observe the behaviour of coral-reef fish. That day, he was watching a grumpy-looking grouper fish as it approached a giant moray eel.

As two of the region's top predators, groupers and morays might be expected to compete for their food and even avoid each other — but Bshary saw them team up to hunt. First, the grouper signalled to the eel with its head, and then the two swam side by side, with the eel dipping into crevices, flushing out fish beyond the grouper's reach and getting a chance to feed alongside. Bshary was astonished by the unexpected

Redouan Bshary
studies social
behaviour of
coral-reef fish.

cooperation; if he hadn't had a snorkel in his mouth, he would have gasped.

This underwater observation was the first in a series of surprising discoveries that Bshary has gone on to make about the social behaviour of fish. Not only can they signal to each other and cooperate across species, but they can also cheat, deceive, console or punish one another — even show concern about their personal reputations. "I have always had a lot of respect for fish," says Bshary. "But one after the other, these behaviours took me by surprise."

His investigations have led him to take a crash course in scuba diving, go beach camping in Egypt and build fake coral reefs in Australia. The work has also destroyed the stereotypical idea that fish are dumb creatures, capable of only the simplest behaviours — and it has presented a challenge to behavioural ecologists in a different field. Scientists who study primates have claimed that human-like behaviours such as cooperation are the sole privilege of animals such as monkeys and apes, and that they helped to drive the evolution of primates' large brains. Bshary — quiet, but afraid of neither adventure nor of contesting others' ideas — has given those scientists reason to think again.

"Redouan has thrown down the gauntlet to us primatologists," says Carel van Schaik, an expert in orang-utan culture at the University of Zurich in Switzerland. "He has made us realize that some of the explanations of primate intelligence that we have cherished don't hold water anymore."

Stream fishing

Bshary says that he was "pre-imprinted to like fish". As a child in Starnberg, Germany, he played constantly in the stream at the edge of the family garden, building dams and pools and trapping fish. Passionate about animal behaviour, he studied evolutionary ecology at the University of Munich, and then did a PhD at the Max Planck Institute for Behavioural Physiology in Starnberg. But for his field work, he journeyed to the Côte d'Ivoire, where he followed tree-living monkeys and discovered that different species collaborate to reduce predator risk.

His PhD supervisor, Ronald Noë, thought it would be "near impossible" to stalk monkeys that leap from tree-top to tree-top, but Bshary seemed to have a flair for it. On occasion, he even camouflaged himself under a leopard skin to imitate one of their predators. And he became fascinated by one question: what makes animals cooperate when standard natural selection would predict selfish behaviour to be the norm?

Noë, a primate behavioural ecologist now at the Hubert Curien Multidisciplinary Institute in Strasbourg, France, had come up with a biological market-based theory of cooperation. It proposed that animals cooperate to trade a specific commodity — such as food — for a service that would promote their survival, such as protection from a predator¹. "An attractive theory — but, there were no strong data to support it," says Bshary.

He looked around for a system where market forces might be operating. And he found one when Hans Fricke, a fish ecologist working at the Max Planck institute, told him the strange tale of barrier-reef fish that operate a remarkable system of cooperation. 'Cleaner' fish, such as the brightly striped wrasse, will nibble parasites off the skin of 'client' fish in small coral territories known as cleaning stations. Bshary realized that this provided a perfect situation in which to test the market theory because client fish seemed to be trading food — in the form of parasites — for a skin-cleaning service. He decided to follow his hunch and study the coral-reef fish.

There was one small problem: Bshary had never been scuba diving. He took his first lessons during a snowy winter in Lake Starnberg, then set off for the Red Sea, setting up camp in Ras Mohammed National Park in Egypt. Together with a few students, Bshary spent two full months a year camped on a scorching beach, sleeping under the stars, eating a

diet of fruit and vegetables and doing four exhausting, 75-minute dives a day. "In the mornings he would wake up and immediately put on his wetsuit and jump straight into the sea," recalls former student Erica van de Waal, now a research fellow at the University of Zurich. Armed with a plastic underwater writing slate, a pencil and a stopwatch, he shadowed client fish, observing their interactions with the wrasse cleaner fish — and soon collected evidence of a well-functioning market. "For me this system was a gold mine," Bshary says, and he mined a lot of gold.

He discovered, for example, that fish did not just trade parasites for skin cleaning; the cleaner fish also cheated on the deal. Rather than eating parasites, they actually preferred the nutritious protective mucus that covers fish skin, and were constantly tempted to take a quick, illicit bite of it. Bshary could count how often this happened — and therefore whether the clients were getting a good or a bad cleaning service — because the clients gave a jolt when they were bitten.

The market theory predicted that if there were lots of clients around, the cleaners would enjoy a seller's market and would risk taking more bites of mucus. This is just like a mechanic getting away with shoddy car services when there are no competing businesses in town. Bshary found this to be true, and he also found that the buyers could protest. Because some client fish roam large territories, they could choose to boycott any cleaning stations that deliver a bad service — just as someone who received a poor car service might travel farther to find a better garage².

While racking up evidence for the market theory, Bshary also observed a range of other social behaviours that had never been seen before in fish. He saw that unsatisfied clients sometimes punish cheating cleaners by chasing them around, and that this punishment makes these fish less likely to cheat³. He saw cleaners ingratiating themselves with certain clients: they gave preference to visiting fish such as groupers, rather than the smaller, local fish that did not have the option of going elsewhere. He found that the cleaners cheated less when they were being watched by other potential clients — a sign that they were buffering their reputations⁴. And he saw reconciliation: if cleaners behaved badly, they then massaged the backs of offended clients with their pelvic fins⁵.

It was all adding up to a catalogue of behaviours worthy of Niccolò Machiavelli's *The Prince* — but it was based on observation alone. Bshary needed to move to an experimental set-up where he could test how the fish behaved. And so in 2003, he began experiments at Lizard Island Research Station on Australia's Great Barrier Reef. He was employed, however, on the other side of the world: first at the University of Liverpool, UK, and now at the University of Neuchâtel, Switzerland. "It was not difficult to sign up to a lifetime of fieldwork at warm coral beach locations," he admits.

Over the next few years Bshary would capture fish in the wild reefs, house them in tanks for the duration of his experiments, then release them. He simulated the choice that cleaners make between parasites and mucus by building moveable plastic plates smeared with prawn, which the fish love, and fish flakes, which they enjoy less. In this set-up, the plates may be snatched away if the cleaners go for the prawns — just like a client fish may swim away if its mucus gets bitten too often. So the cleaners learned to cooperate and eat fish flakes instead.

Such experiments take patience: some fish take a month just to adjust to the tanks. But in this way, Bshary proved that all the behaviours he had observed in the wild could be repeated under experimental conditions. And he discovered even more bizarre facts about the social lives of fish. In one experiment, he showed that when cleaners work in male and female pairs, as frequently happens in the wild, they are much less likely to cheat than when they work alone⁶; and that this is mostly because the female gets punished by being chased around by the male if she slacks off⁷.



A fake moray, made from a laminated photo, fools a grouper fish.

Perhaps the laboratory's most imaginative experiment involved the construction of an entire fake coral reef, complete with dummy eels. The job fell to doctoral student Alex Vail, who glued together bits of coral rubble and put them in a hot-tub-sized tank. Vail then made models of moray eels by printing, gluing together and laminating two life-sized photographs, and attaching nylon strings that allowed him to pull the fake eels out of the fake coral, like a puppet. (Vail subsequently went on to a successful career in underwater filming.) Using this set-up, the team explored the behaviour that so shocked Bshary when he observed it in 1998: a grouper and eel teaming up to flush out fish to eat. They showed that the grouper quickly learnt to signal — by turning and shaking its head — only to those moray eels that responded by moving towards, rather than away from, the fake reef.

Bshary amassed ample evidence that fish engage in a range of social behaviours, and he assumed that all of them resulted from simple evolution at work. Natural selection favoured fish that could learn, by simple association, which choices allowed them to efficiently rid themselves of parasites or access food.

By 2010, Bshary's thoughts were turning back to the world of primatology, in which he had been immersed during his PhD. He knew that he had observed in fish many of the behaviours that primatologists had shown in monkeys and apes. But primatologists had made grander claims for their observations. The 'social brain' theory argues that primates evolved brains that are large for their body size to manage their unusually complex social systems. Only primate brains, the theory says, have the depth of cognitive analysis necessary to cooperate, deceive and solve other problems in a social world.

Bshary disagreed. Maybe, he thought, these particular social behaviours in primates were also learnt by simple association and did not require the extra computing power of their big brains. And his findings meshed with those emerging from studies on the social behaviours of other animals, ranging from elephants to birds. "I think primatologists tend to make big claims because they look up the evolutionary chain and compare the primates' behaviours to humans, instead of looking down the evolutionary chain to see if the phenomena also existed in lower species," he says.

At the time, primatologists were certainly not looking at fish. But that changed when Bshary teamed up with primatologist Sarah Brosnan at Georgia State University in Atlanta to directly pit the skills of cleaner fish against capuchin monkeys, chimpanzees and orang-utans in a foraging test. Each animal was presented with food on two differently coloured plates, one of which was a permanent fixture in their tanks or pens, whereas the other was temporary. The challenge was to learn to eat

from the temporary plate first, before it disappeared — and the scientists counted how many trials it took for the animal to figure this out.

The cleaner fish solved the problem first⁸; they have evolved in their ecological niche to preferentially feast from visiting clients before they disappear. For fun, Bshary set up an equivalent 'foraging' test for his four-year-old daughter, complete with temporary and permanent plates, each bearing one chocolate M&M. In a series of 100 different trials, she never learnt to eat from the temporary plate.

The fish, meanwhile, were already aceing a more advanced test. When Bshary and Brosnan switched the coloured plates so that the permanent one suddenly became temporary and vice versa, the fish again understood the switch faster than the apes did (and equally as fast as the capuchins)⁸. This is known as reversal learning — and when the primatologists read that result, they took note. "Reversal learning has often been touted as the gold standard of general cognitive abilities," says van Schaik — a sophisticated skill that correlates with brain size. "Since small-brained fish do it quite well, maybe we'll have to abandon this idea."

"The ball is in our court," says evolutionary psychologist Robin Dunbar of the University of Oxford, UK, who developed the social brain theory. Dunbar now accepts that the evolution of large brains was not driven by the need to carry out single 'smart' behaviours such as cooperation or deception. But that doesn't mean the social brain theory has to be abandoned, he says — just refined. He and other primatologists now propose that primates evolved bigger brains because they needed an all-round high level of general intelligence to survive the pressures of living in tight social groups — for example, to recognize large numbers of individuals and remember their complicated genetic and hierarchical relationships.

Fish, which tend to have one-on-one interactions and live in loose schools, do not need to multi-task in quite the same way, Dunbar says. "It may boil down to the speed of cognitive processing and accuracy of judgement," he suggests.

Intelligence tests

Michael Tomasello, an evolutionary psychologist at the Max Planck Institute for Evolutionary Anthropology in Leipzig, Germany, bounces the ball right back to Bshary, challenging him to show how smart fish really are. "Perhaps the most pressing question is how flexible and general fish cognition is," he says — something Bshary is already testing by designing further fish intelligence tests.

The mysteries of the fish brain deepened in 2009, when Bshary's team chanced across a habitat in the reefs around Lizard Island that had relatively few fish and therefore less competition and social complexity. To Bshary's surprise, the cleaner fish there turned out to be much less socially smart than cleaner fish just 20 metres away⁹. But their skill level may be optimal for their environment — another hypothesis that he now plans to explore.

Whatever the next instalment brings, colleagues say that Bshary has already shifted a view of animal cognition in which humans and their primate cousins tower over everything else. "Primate chauvinism may now be poised to decline, thanks in large part to Bshary's fish work," says primatologist and ethologist Frans de Waal of Emory University in Atlanta, Georgia. "They now really do have to take on board that most species are going to have a type of smart intelligence." ■

Alison Abbott is Nature's senior European correspondent.

- Noë, R. & Hammerstein, P. *Trends Ecol. Evol.* **10**, 336–339 (1995).
- Bshary, R. & D'Souza, A. in *Animal Communication Networks* (ed. McGregor, P. K.) 521–539 (Cambridge Univ. Press, 2005).
- Bshary, R. & Grutter, A. S. *Biol. Lett.* **1**, 396–399 (2005).
- Bshary, R. & Grutter, A. S. *Nature* **441**, 975–978 (2006).
- Bshary, R. & Würth, M. *Proc. R. Soc. Lond. B* **268**, 1495–1501 (2001).
- Bshary, R., Grutter, A. S., Willener, A. S. T. & Leimar, O. *Nature* **455**, 964–966 (2008).
- Raihani, N. J., Pinto, A. I., Grutter, A. S., Wismer, S. & Bshary, R. *Proc. R. Soc. Lond. B* **279**, 365–370 (2012).
- Salwiczek, L. H. *et al.* *PLoS ONE* **7**, e49068 (2012).
- Wismer, S., Pinto, A. I., Vail, A. L., Grutter, A. S. & Bshary, R. *Ethology* **120**, 519–531 (2014).

COMMENT

ECONOMICS New metric captures accumulation of productive information **p.420**

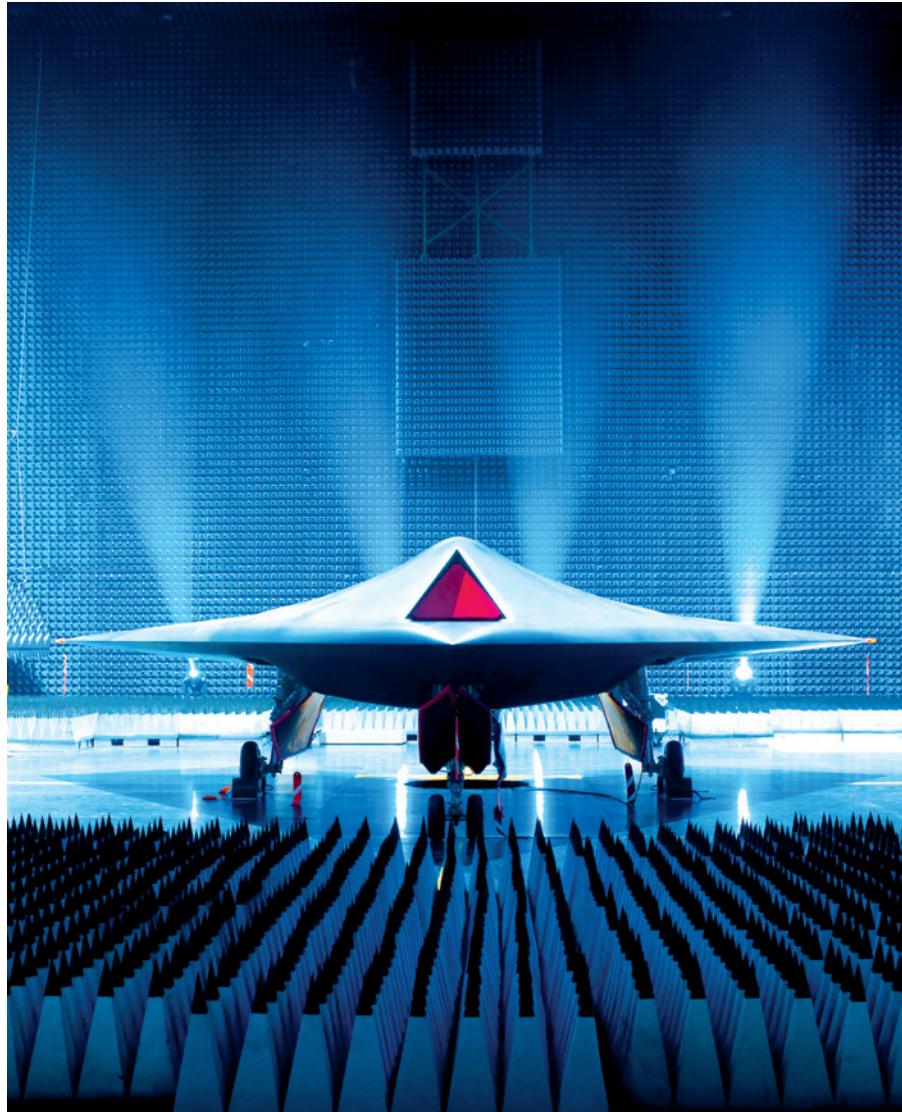
CHEMISTRY Tracing the evolution of the lab, from furnace to fume hood **p.422**

CONSERVATION Deforestation soaring in the Amazon, satellite data show **p.423**

INSTRUMENTS Could microscope found in mud be an original Leeuwenhoek? **p.423**



TONY GARNER/BAE



BAE Systems' Taranis drone has autonomous elements, but relies on humans for combat decisions.

Ethics of artificial intelligence

Four leading researchers share their concerns and solutions for reducing societal risks from intelligent machines.

STUART RUSSELL Take a stand on AI weapons

*Professor of computer science,
University of California, Berkeley*

The artificial intelligence (AI) and robotics communities face an important ethical decision: whether to support or oppose the development of lethal autonomous weapons systems (LAWS).

Technologies have reached a point at which the deployment of such systems is — practically if not legally — feasible within years, not decades. The stakes are high: LAWS have been described as the third revolution in warfare, after gunpowder and nuclear arms.

Autonomous weapons systems select and engage targets without human intervention; they become lethal when those targets include humans. LAWS might include, for example, armed quadcopters that can search for and eliminate enemy combatants in a city, but do not include cruise missiles or remotely piloted drones for which humans make all targeting decisions.

Existing AI and robotics components can provide physical platforms, perception, motor control, navigation, mapping, tactical decision-making and long-term planning. They just need to be combined. For example, the technology already demonstrated for self-driving cars, together with the human-like tactical control learned by DeepMind's DQN system, could support urban search-and-destroy missions.

Two US Defense Advanced Research Projects Agency (DARPA) programmes foreshadow planned uses of LAWS: Fast Lightweight Autonomy (FLA) and Collaborative Operations in Denied Environment (CODE). The FLA project will program tiny rotorcraft to manoeuvre unaided at high speed in urban areas and inside buildings. CODE aims to develop teams of autonomous aerial vehicles carrying out "all steps of a strike mission — find, fix, track, target, engage, assess" in situations in which enemy signal-jamming makes communication with a human commander impossible. Other ▶

► countries may be pursuing clandestine programmes with similar goals.

International humanitarian law — which governs attacks on humans in times of war — has no specific provisions for such autonomy, but may still be applicable. The 1949 Geneva Convention on humane conduct in war requires any attack to satisfy three criteria: military necessity; discrimination between combatants and non-combatants; and proportionality between the value of the military objective and the potential for collateral damage. (Also relevant is the Martens Clause, added in 1977, which bans weapons that violate the “principles of humanity and the dictates of public conscience.”) These are subjective judgments that are difficult or impossible for current AI systems to satisfy.

The United Nations has held a series of meetings on LAWS under the auspices of the Convention on Certain Conventional Weapons (CCW) in Geneva, Switzerland. Within a few years, the process could result in an international treaty limiting or banning autonomous weapons, as happened with blinding laser weapons in 1995; or it could leave in place the status quo, leading inevitably to an arms race.

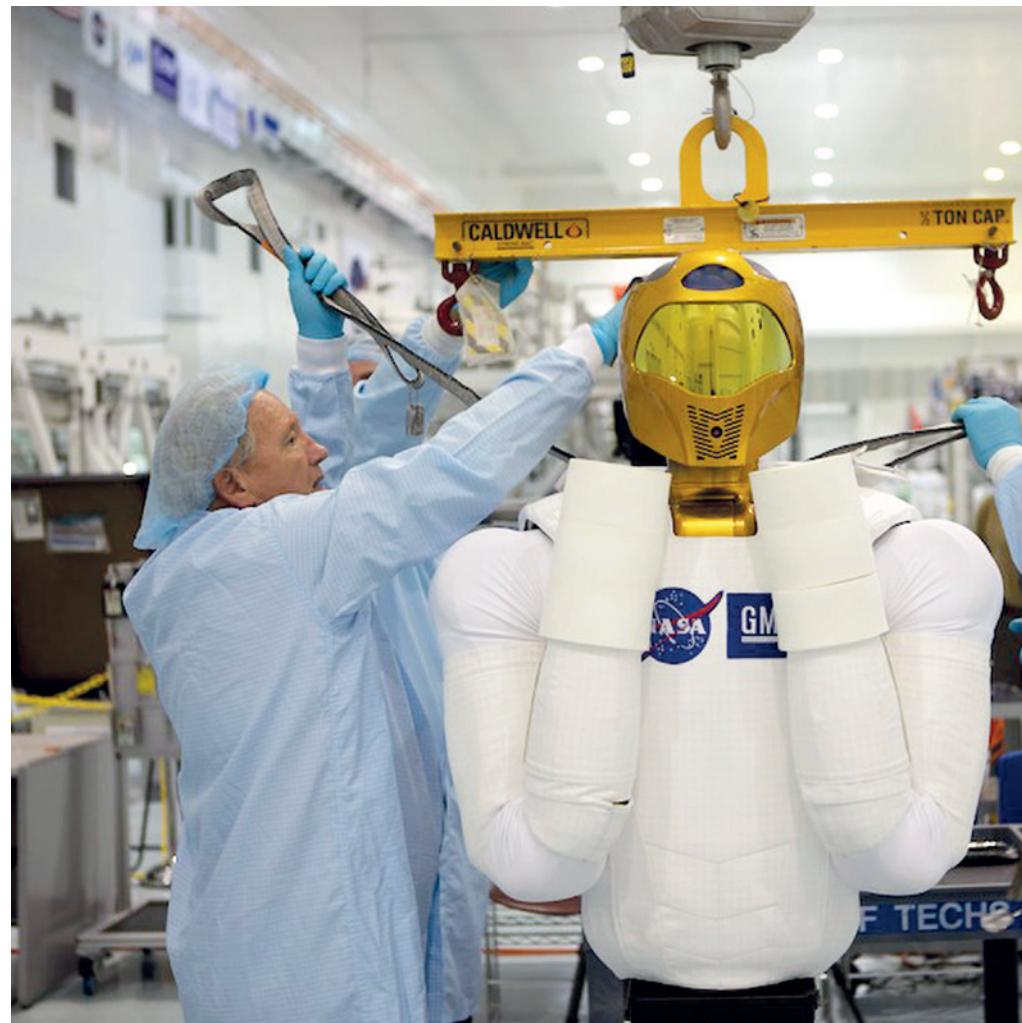
As an AI specialist, I was asked to provide expert testimony for the third major meeting under the CCW, held in April, and heard the statements made by nations and non-governmental organizations. Several countries pressed for an immediate ban. Germany said that it “will not accept that the decision over life and death is taken solely by an autonomous system”; Japan stated that it “has no plan to develop robots with humans out of the loop, which may be capable of committing murder” ([see go.nature.com/fwrc1](http://go.nature.com/fwrc1)).

The United States, the United Kingdom and Israel — the three countries leading the development of LAWS technology — suggested that a treaty is unnecessary because they already have internal weapons review processes that ensure compliance with international law.

Almost all states who are party to the CCW agree with the need for ‘meaningful human control’ over the targeting and engagement decisions made by robotic weapons. Unfortunately, the meaning of ‘meaningful’ is still to be determined.

The debate has many facets. Some argue that the superior effectiveness and selectivity of autonomous weapons can minimize civilian casualties by targeting only combatants. Others insist that LAWS will lower the threshold for going to war by making it possible to attack an enemy while incurring no immediate risk; or that they will enable terrorists and non-state-aligned combatants to inflict catastrophic damage on civilian populations.

LAWS could violate fundamental principles of human dignity by allowing machines



NASA's Robonaut 2 could be used in medicine and industry as well as space-station construction.

to choose whom to kill — for example, they might be tasked to eliminate anyone exhibiting ‘threatening behaviour’. The potential for LAWS technologies to bleed over into peacetime policing functions is evident to human-rights organizations and drone manufacturers.

In my view, the overriding concern should be the probable endpoint of this technological trajectory. The capabilities of autonomous weapons will be limited more by the laws of physics — for example, by constraints on range, speed and payload — than by any deficiencies in the AI systems that control them. For instance, as flying robots become smaller, their manoeuvrability increases and their ability to be targeted decreases. They have a shorter range, yet they must be large enough to carry a lethal payload — perhaps a one-gram shaped charge to puncture the human cranium. Despite the limits imposed by physics, one can expect platforms deployed in the millions, the agility and lethality of which will leave humans utterly defenceless. This is not a desirable future.

The AI and robotics science communities,

represented by their professional societies, are obliged to take a position, just as physicists have done on the use of nuclear weapons, chemists on the use of chemical agents and biologists on the use of disease agents in warfare. Debates should be organized at scientific meetings; arguments studied by ethics committees; position papers written for society publications; and votes taken by society members. Doing nothing is a vote in favour of continued development and deployment.

SABINE HAUERT
**Shape the debate,
don't shy from it**

Lecturer in robotics, University of Bristol

Irritated by hyped headlines that foster fear or overinflate expectations of robotics and artificial intelligence (AI), some researchers have stopped communicating with the



media or the public altogether.

But we must not disengage. The public includes taxpayers, policy-makers, investors and those who could benefit from the technology. They hear a mostly one-sided discussion that leaves them worried that robots will take their jobs, fearful that AI poses an existential threat, and wondering whether laws should be passed to keep hypothetical technology ‘under control’. My colleagues and I spend dinner parties explaining that we are not evil but instead have been working for years to develop systems that could help the elderly, improve health care, make jobs safer and more efficient, and allow us to explore space or beneath the oceans.

Experts need to become the messengers. Through social media, researchers have a public platform that they should use to drive a balanced discussion. We can talk about the latest developments and limitations, provide the big picture and demystify the technology. I have used social media to crowd-source designs for swarming nanobots to treat cancer. And I found my first PhD student through his nanomedicine blog.

The AI and robotics communities need

thought leaders who can engage with prominent commentators, such as physicist Stephen Hawking and entrepreneur-inventor Elon Musk, and set the agenda at international meetings such as the World Economic Forum in Davos, Switzerland. Public engagement also drives funding. Crowdfunding for JIBO, a personal robot for the home developed by Cynthia Breazeal, at the Massachusetts Institute of Technology (MIT) in Cambridge, raised more than US\$2.2 million.

There are hurdles. First, many researchers have never tweeted, blogged or made a YouTube video. Second, outreach is ‘yet another thing to do’, and time is limited. Third, it can take years to build a social-media following that makes the effort worthwhile. And fourth, engagement work is rarely valued in research assessments, or regarded seriously by tenure committees.

Training, support and incentives are needed. All three are provided by Robohub.org, of which I am co-founder and president. Launched in 2012, Robohub is dedicated to connecting the robotics community to the public. We provide crash courses in science communication at major AI and robotics conferences on how to use social media efficiently and effectively. We invite professional science communicators and journalists to help researchers to prepare an article about their work. The communicators explain how to shape messages to make them clear and concise and avoid pitfalls, but we make sure the researcher drives the story and controls the end result. We also bring video cameras and ask researchers who are presenting at conferences to pitch their work to the public in five minutes. The results are uploaded to YouTube. We have built a portal for disseminating blogs and tweets, amplifying their reach to tens of thousands of followers.

I can list all the benefits of science communication, but the incentive must come from funding agencies and institutes. Citations cannot be the only measure of success for grants and academic progression; we must also value shares, views, comments and likes. MIT robotics researcher Rodney Brooks’s classic 1986 paper on the ‘subsumption architecture’, a

NATURE.COM
For more, see the
Nature Insight on
machine intelligence:
go.nature.com/eizewc

bio-inspired way to program robots to react to their environment, gathered nearly 10,000 citations in three

decades (R. Brooks *IEEE J. Robot. Automat.* **2**, 14–23; 1986). A video of Sawyer, a robot developed by Brooks’s company Rethink Robotics, received more than 60,000 views in one month (see go.nature.com/jqwfmz). Which has had more impact on today’s public discourse?

Governments, research institutes, business-development agencies, and research and industry associations do welcome and fund outreach and science-communication efforts. But each project develops its own strategy, resulting in pockets of communication that have little reach.

In my view, AI and robotics stakeholders worldwide should pool a small portion of their budgets (say 0.1%) to bring together these disjointed communications and enable the field to speak more loudly. Special-interest groups, such as the Small Unmanned Aerial Vehicles Coalition that is promoting a US market for commercial drones, are pushing the interests of major corporations to regulators. There are few concerted efforts to promote robotics and AI research in the public sphere. This balance is badly needed.

A common communications strategy will empower a new generation of robotists that is deeply connected to the public and able to hold its own in discussions. This is essential if we are to counter media hype and prevent misconceptions from driving perception, policy and funding decisions.

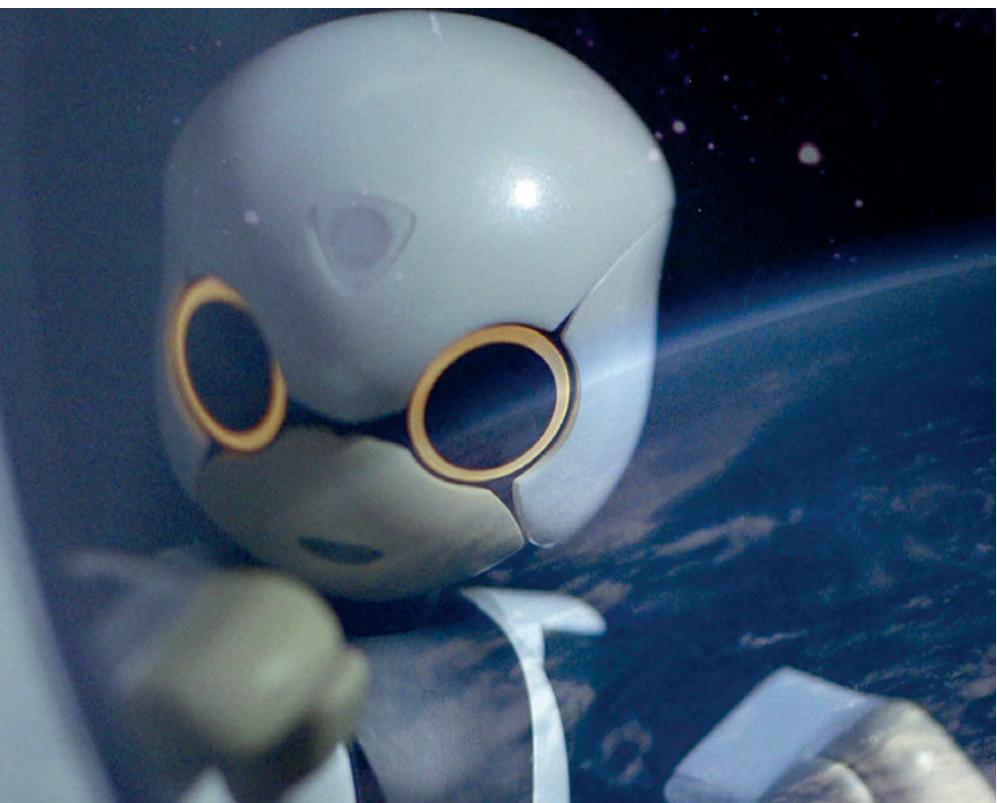
RUSS ALTMAN

Distribute AI benefits fairly

Professor of bioengineering, genetics, medicine and computer science,
Stanford University

Artificial intelligence (AI) has astounding potential to accelerate scientific discovery in biology and medicine, and to transform health care. AI systems promise to help make sense of several new types of data: measurements from the ‘omics’ such as genomics, proteomics and metabolomics; electronic health records; and digital-sensor monitoring of health signs.

Clustering analyses can define new syndromes — separating diseases that were thought to be the same and unifying others that have the same underlying defects. Pattern-recognition technologies may match disease states to optimal treatments. For example, my colleagues and I are identifying groups of patients who are likely to respond to drugs that regulate the immune system on the basis of clinical and transcriptomic features.



Kirobo, Japan's first robot astronaut, was deployed to the International Space Station in 2013.

In consultations, physicians might be able to display data from a 'virtual cohort' of patients who are similar to the one sitting next to them and use it to weigh up diagnoses, treatment options and the statistics of outcomes. They could make medical decisions interactively with such a system or use simulations to predict outcomes on the basis of the patient's data and that of the virtual cohort.

I have two concerns. First, AI technologies could exacerbate existing health-care disparities and create new ones unless they are implemented in a way that allows all patients to benefit. In the United States, for example, people without jobs experience diverse levels of care. A two-tiered system in which only special groups or those who can pay — and not the poor — receive the benefits of advanced decision-making systems would be unjust and unfair. It is the joint responsibility of the government and those who develop the technology and support the research to ensure that AI technologies are distributed equally.

"AI technologies could exacerbate existing health-care disparities and create new ones."

Second, I worry about clinicians' ability to understand and explain the output of high-performance AI systems. Most health-care providers will not accept a

complex treatment recommendation from a decision-support system without a clear description of how and why it was reached.

Unfortunately, the better the AI system, the harder it often is to explain. The features that contribute to probability-based assessments such as Bayesian analyses are straightforward to present; deep-learning networks, less so.

AI researchers who create the infrastructure and technical capabilities for these systems need to engage doctors, nurses, patients and others to understand how they will be used, and used fairly.

MANUELA VELOSO Embrace a robot-human world

*Professor of computer science,
Carnegie Mellon University*

Humans seamlessly integrate perception, cognition and action. We use our sensors to assess the state of the world, our brains to think and choose actions to achieve objectives, and our bodies to execute those actions. My research team is trying to build robots that are capable of doing the same — with artificial sensors (cameras, microphones and scanners), algorithms and actuators, which control the mechanisms.

But autonomous robots and humans differ greatly in their abilities. Robots may always have perceptual, cognitive and actuation limitations. They might not be able to fully perceive a scene, recognize or manipulate any object, understand all spoken or written language, or navigate in any terrain. I think that robots will complement humans, not supplant them. But robots need to know when to ask for help and how to express their inner workings.

To learn more about how robots and humans work together, for the past three years we have shared our laboratory and buildings with four collaborative robots, or CoBots, which we developed. The robots look a bit like mechanical lecterns. They have omnidirectional wheels that enable them to steer smoothly around obstacles; camera and lidar systems to provide depth vision; computers for processing; screens for communication; and a basket to carry things in.

Early on, we realized how challenging real environments are for robots. The CoBots cannot recognize every object they encounter; lacking arms or hands they struggle to open doors, pick things up or manipulate them. Although they can use speech to communicate, they may not recognize or understand the meaning of words spoken in response.

We introduced the concept of 'symbiotic autonomy' to enable robots to ask for help from humans or from the Internet. Now, robots and humans in our building aid one another in overcoming the limitations of each other.

CoBots escort visitors through the building or carry objects between locations, gathering useful information along the way. For example, they can generate accurate maps of spaces, showing temperature, humidity, noise and light levels, or WiFi signal strength. We help the robots to open doors, press lift buttons, pick up objects and follow dialogue by giving clarifications.

There are still hurdles to overcome to enable robots and humans to co-exist safely and productively. My team is researching how people and robots can communicate more easily through language and gestures, and how robots and people can better match their representations of objects, tasks and goals.

We are also studying how robot appearance enhances interactions, in particular how indicator lights could reveal more of a robot's inner state to humans. For instance, if the robot is busy, its lights may be yellow, but when it is available they are green.

Although we have a way to go, I believe that the future will be a positive one if humans and robots can help and complement each other. ■ SEE INSIGHT P.435



IVAN ALVARADO/REUTERS/CORBIS

Lithium is produced from brine in Chile's Atacama Desert.

INFORMATION THEORY

Knowledge and know-how

Philip Ball examines a study unpicking the broad ramifications of information flows.

The Atacama desert in Chile is the world's biggest single source of lithium. But Chile does not make its own lithium batteries; for that, the metal is exported to countries as distant as South Korea and China. Why?

"Our world is marked by great international differences in countries' ability to crystallize imagination," writes César Hidalgo in *Why Information Grows*: that is, in the ability to turn the inventiveness of the human mind into concrete, saleable products that embody information.

Hidalgo's book is largely concerned with explaining why those differences exist, and what the economic consequences are. In doing so, Hidalgo, a Chilean physicist at the Media Lab of the Massachusetts Institute of Technology in Cambridge, aims to embed his ideas in a broad view of how information economies work both in society and in nature, from Silicon Valley entrepreneurship to genetics.

Information can now be accessed so easily



Why Information Grows: The Evolution of Order, from Atoms to Economies
CÉSAR HIDALGO
Basic: 2015

— you can find out about pretty much anything from anywhere, or so it sometimes seems — that it might naively be expected to reduce discrepancies in the local know-how needed to "crystallize imagination". But differences in the character and performance of national economies persist in ways that are not simply the result of inertia; nor can they be explained solely by different labour costs. The core of *Why Information Grows* is about the acquisition (or not) of the capability to put information to use. Many of Chile's recent exports to South Korea were copper (another resource that Chile has in abundance). The reverse trade was largely in vehicles and their parts. The trade balance

lies in Chile's favour financially, but Hidalgo points out that this masks a deficit in the flow of embodied information. That does not, of course, mean that Chilean automotive engineers are ignorant. Rather, the know-how to make cars reflects a whole suite of considerations about capacity: a trained workforce, access to materials, distribution channels, reputation, trading partners and so on. To truly understand these macroeconomic transactions, Hidalgo argues, we need to delve into the infrastructure and networks that make knowledge productive and allow know-how to accumulate.

Part of this understanding has already been developed by economists, in particular Ronald Coase, whose 1930s work on transaction costs and the origin and growth of firms won him the 1991 economics Nobel. More recently, social capital — connections, and the trust they engender — has been recognized

NATURE.COM
For more on science in culture, see:
nature.com/booksandarts

as a key factor in the health of economies. Hidalgo's particular contribution here is to weave these ideas into a theory that provides a conceptual picture of a society's collective accumulation of productive information "in firms and networks of firms". He argues that these entities are not so much the result of that accumulation as the means by which our societies have evolved to make it possible.

What emerges is a new measure of economic complexity that acknowledges not just gross domestic product, but also diversity in the types of product that a nation produces. This index, Hidalgo says, offers a metric that predicts long-term economic growth. His thesis here is dense and not always lucidly explained, but it contains some innovative thinking about what drives growth that could help us to navigate the turbulence of the ever more interconnected global economy.

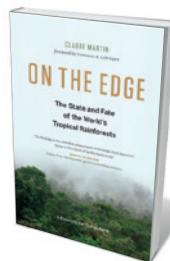
Hidalgo wants to take his ideas about information and its productive exploitation beyond economics, especially into biology and ecology. In a nutshell, the task is to restore notions of meaning to information theory. The seminal text *The Mathematical Theory of Communication* (Univ. Illinois Press, 1949) by Claude Shannon and Warren Weaver explicitly expunged the idea that information in itself has meaning, as Hidalgo reminds us. But we often persist in thinking otherwise, not least in the delusions that we can do science by mining big data without a guiding theory, or that genomes contain 'instructions' for making an organism.

The distinction between knowledge and know-how, which seems evident enough in economics, can be profitably made in the natural sciences too. A gene sequence is information; know-how refers to the ability to make something from it, and this cannot reside in the information itself. "DNA has no knowhow and cannot unpack itself; it is a slave to the machinery needed to unpack it," says Hidalgo. This seems obvious, but it is worth restating, if only to remind us that this unpacking process (and not further sequencing per se) is now the most important unanswered question in genomics.

Whether all this really justifies Hidalgo's claim to explain "what information is, where it comes from, and why it grows" is another matter. For one thing, with quantum mechanics being reframed as an information theory, it is unclear whether any classical arguments can give the whole picture. But that does not detract from the stimulating new perspectives on offer. Hidalgo has identified a fertile seam, and all his book really needs (apart from some prose-tightening) is for the title to be turned into a question. ■

Philip Ball is a writer based in London. His latest book is *Invisible*.
e-mail: p.ball@btinternet.com

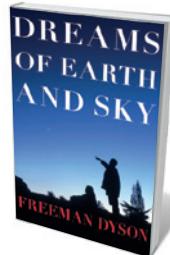
Books in brief



On the Edge: The State and Fate of the World's Tropical Rainforests

Claude Martin GREYSTONE (2015)

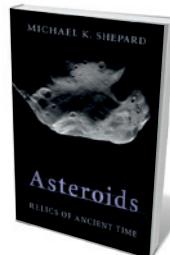
However crucial in terms of biodiversity, climate moderation and the carbon cycle, rainforests are at bay. Almost one-fifth of the Amazon, for instance, has been lost over 50 years. In this definitive report to think tank the Club of Rome, conservation veteran Claude Martin and contributors cover the territory from the bad old days of error-ridden global forest monitoring to today's smart mapping and strategies for sustainable management. Martin predicts a further 100-million-hectare loss of primary forest by 2050 — but delivers 17 steps for averting the worst. Key reading for environmental policy-makers.



Dreams of Earth and Sky

Freeman Dyson NEW YORK REVIEW BOOKS (2015)

"The big blunders in this book are not accidental but intentional." So physicist Freeman Dyson proclaims of the maverick ideas percolating through this collection of his *New York Review of Books* essays and book reviews from 2006 to 2014. Dyson delivers his iconoclasm with humour, honesty and off-the-cuff brilliance, whether riffing off microbiologist Carl Woese's concept of living systems by seeing both thunderstorms and butterflies as "patterns of organization", imagining ecology enriched by genome designers, or mulling over how Mary Shelley's *Frankenstein* cast science as "an agent of doom".



Asteroids: Relics of Ancient Time

Michael K. Shepard CAMBRIDGE UNIVERSITY PRESS (2015)

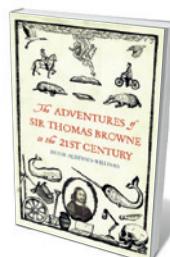
The chunks of broken protoplanet we call asteroids are potential killers as well as scientific gold mines. Both thrills are explored in this comprehensive study of these "little worlds". Planetary geoscientist Michael Shepard recounts the history of asteroid spotting, from pioneers such as Giuseppe Piazzi — who detected the first, Ceres, in 1801 — to today's Catalina Sky Survey, run by the University of Arizona. Interweaving anecdotes from his own work with copious technical detail, Shepard leads us on an expert tour of a fecund branch of astronomy and its lessons for planetary science.



Breathing Space: The Natural and Unnatural History of Air

Mark Everard ZED (2015)

The air — all 51 quadrillion cubic kilometres of it — is the biggest ecosystem on Earth, an atmospheric ocean connecting all life. Yet as science writer Mark Everard shows in his original study, the importance of this layered envelope of gases and aerosols is hardly reflected in policy. He argues for reassessment, looking in turn at the nexus of atmosphere, biosphere and human culture, and abuses such as air pollution and climate change. To secure "breathing space", Everard calls for an approach integrating management of the atmosphere with that of other Earth systems.



The Adventures of Sir Thomas Browne in the 21st Century

Hugh Aldersey-Williams GRANTA (2015)

A playful erudition permeates this biography of Thomas Browne, the seventeenth-century English polymath who wrote the debunking classic *Pseudodoxia Epidemica*, coined 784 new words and had an eye for patterns in nature. Hugh Aldersey-Williams zips between our time and Browne's, and through the East Anglian landscapes Browne knew, to reveal the man and his work. Browne emerges as an exemplar of synthesized knowledge — and as such, curiously at home in today's cultural remix of science and humanities. **Barbara Kiser**

LABORATORY HISTORY

The chemistry chronicles

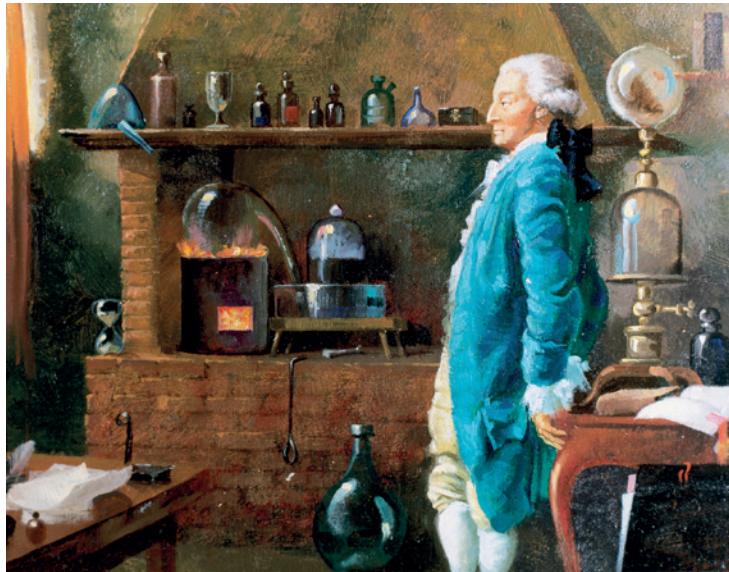
Derek Lowe relishes a celebration of the lab on its long march through time.

Ask non-scientists to picture a laboratory, and the odds are that they will imagine the kind of space that a contemporary chemist would work in, complete with lab bench, fume cupboard, sink and shelves of glassware. Chemistry was probably the first of the sciences to get a room of its own, and in *The Matter Factory*, Peter Morris, keeper of research projects at London's Science Museum, offers the first book-length treatment of how this happened and what has changed in labs over the years.

Laboratory history is inseparable from the history of chemistry. It takes us from scientific instrumentation to teaching, from the discipline's beginnings as the hidden art of alchemy to its modern status as a required course in science. It also takes in the rise of industrial research, the gradual raising of consciousness about safety and the personal impact of famous chemists. No discipline is immune to fashion, and a luminary — such as Robert Bunsen, who co-discovered spectroscopic analysis, or organic-chemistry pioneer Justus von Liebig — could set the tone for years by building a laboratory to his personal specifications.

As Morris shows, over all this stretches culture. Chemistry moved from a broadly French and English occupation during the Enlightenment, an era of experiments on gases and chemical compositions, to a German one in the mid-nineteenth century, with the Americans eventually beginning to take notes and draw up plans of their own in the twentieth century. Each nation had its own style, which blended with the practical aspects of a workplace to create distinctive looks. The most useful features (such as benches with drawers, and dedicated lines for gases and steam) are still to be found today.

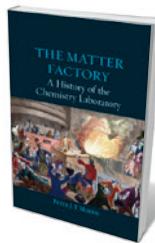
Morris deftly keeps all these threads from getting too tangled. *The Matter Factory* starts in the alchemist's lair of the medieval era, dominated by the largest, hottest furnace available. The book makes clear that many engravings and paintings of alchemists at work must be inaccurate, because they were drawn by people suspicious of



Antoine Lavoisier was central to the eighteenth-century chemical revolution.

the whole enterprise. (Alchemists' own drawings, such as those in the seventeenth-century text *Mutus Liber*, tended toward the wildly allegorical and the wilfully obscure.) The German scientist Georgius Agricola's 1556 *De Re Metallica* (*On the Nature of Metals*) is probably the first reliable guide to early lab technique such as the handling of strong acids. In later illustrations, the furnace shrinks, then disappears entirely, and tables and benches appear.

Fume hoods began to take on a modern look by the 1920s, but separate cupboard-like spaces for experimentation (at first unventilated) go back to at least the mid-nineteenth century. Photographs from that time show benches and shelving progressively stretching across the floors and walls, and lines for water, gas, steam and (by the early twentieth century) electricity threading into the picture. Teaching labs gradually separate themselves from research labs, and industrial labs begin to stand on their own. Along the way, we chemists lose such amenities as the chemical museum, which displayed interesting specimens, compounds and lab equipment — a once common annex to



The Matter Factory: A History of the Chemistry Laboratory
PETER J.T. MORRIS
Reaktion: 2015.

many large laboratories, now completely forgotten.

The nineteenth century saw perhaps the greatest number of changes in laboratory layout, as new instruments such as Bunsen burners and new styles of working, including team research, proliferated. Much of the book focuses on this era. The pace speeds up noticeably in the twentieth century, probably because the major features of the modern lab were already largely in place, down to the pegboard over the sink for dry-ing glassware.

Today, one chemistry lab tends to look very much like another. In my own career, I have 'lived' in at least 14 labs, and their similarities are much more numerous than their differences, which largely come down to how instruments such as high-performance-liquid-chromatography equipment are arranged around the walls. If a chemist from, say, the early nineteenth century entered one, the glassware would be familiar, as would the bench and fume cupboard, although there might be a few cries of "Luxury!" at the sight of disposable pipettes and electronic balances. (The humming boxes around the periphery would, of course, be harder to place.)

The changes to labs of the future will probably come down to variations in the number and capacity of automated instruments. But the space will probably look broadly similar to what we have now, which will no doubt disappoint some industrial designers looking to make a big splash. *The Matter Factory*, however, is the story of the years (and centuries) during which such splashes could be made, when chemistry was finding out what it could do. It covers a lot of ground, and brings together many old drawings, plans and photographs that are otherwise scattered through a bewildering literature trail. It should remain the definitive history of the chemistry lab for many years. ■

Derek Lowe has been working in drug discovery labs at several companies since 1989, and blogs at *In the Pipeline* (www.pipeline.corante.com).
e-mail: derekb.lowe@gmail.com

Correspondence

The mystery of the microscope in mud

A rare brass microscope was discovered in mud dredged from canals in Delft, the Netherlands, in December 2014. The instrument (pictured, left) is thought to have been made by Dutch pioneer microscopist Antony van Leeuwenhoek (1632–1723; see P. Ball *Nature* **520**, 156; 2015) and is an important find — if it is genuine. The last Leeuwenhoek microscope in public hands was a silver instrument that sold for more than US\$500,000 in 2009.

The ten microscopes attributed to Leeuwenhoek are diminutive, comprising a biconvex lens sandwiched between two riveted body plates. A positioning screw holds a stage block and a pin to secure the specimen. A smaller screw moves the stage block to focus the image.

The Delft microscope's design and dimensions are comparable to those of a Leeuwenhoek microscope in the Boerhaave Museum in Leiden. Although the lens is abraded, it would have given about 160 \times magnification.

The Delft find cannot be a replica. First, it has a rounded body plate — a comparable microscope in the Boerhaave Museum is rectangular (pictured, right), and a copy would have an identical configuration. Second, my analysis of the screw thread at the Cavendish Laboratory at the University of Cambridge, UK, shows that it is unlike threads produced by modern mechanical methods. The distance between the threads (or pitch) of the main screw is 0.9 millimetres, comparable with that in authentic instruments.

I am preparing a protocol through which the production details of all existing Leeuwenhoek microscopes can be scrutinized. Scanning electron microscopy should then be able to detect any forgeries.

Brian J. Ford *Cardiff University, UK.*
brianjford@cardiff.ac.uk



Deforestation soars in the Amazon

Deforestation of the Brazilian Amazon is on the rise and not declining as you imply (*Nature* **520**, 20–23; 2015). Forest clearance has exploded since August 2014, according to satellite data from Brazil's National Institute for Space Research and the sustainable-development organization Imazon. Monthly clearance rates are several times those for the same months in the preceding two years.

The data on deforestation drivers in the studies you discuss predate the recent deforestation surge (D. Nepstad *et al.* *Science* **344**, 1118–1123; 2014 and H. K. Gibbs *et al.* *Science* **347**, 377–378; 2015). Brazil's currency, the real, is in free-fall, with no end in sight. Exchange rates against the US dollar mean that exports from Brazil are now considerably more profitable than last year, when the deforestation decline was occurring.

Prospects for the near term are not good. A key 2008 resolution from the Central Bank of Brazil, which links financing from public banks to environmental compliance, now offers only fragile protection in the face of a congress dominated by those in favour of agribusiness. Creation of new protected areas has been essentially paralysed, existing reserves continue to have their official status

removed, and government expenditure on enforcing environmental laws has been cut by 72%. Furthermore, political appointments are sending an anti-environmental signal to deforesters (see *Nature* **517**, 251–252; 2015), and plans for Amazonian roads continue as fast as money allows. The "battle for the Amazon" is far from won.

Philip M. Fearnside *National Institute for Research in Amazonia (INPA), Manaus, Brazil.*
pmfearn@inpa.gov.br

UK bioscientists push for crop policy

In assessing the next generation of plant-breeding techniques such as gene editing (*Nature* **520**, 131–132; 2015), the European Union (EU) should take into account criticisms of its current regulatory system for new crop varieties.

Last year's statement by the UK Biotechnology and Biological Sciences Research Council on genetic techniques for crop improvement suggests that EU regulations need to move beyond production methods and focus more on the benefits of newly introduced genetic traits (see go.nature.com/gsoiza).

The UK National Institutes of Bioscience, which contributed to the statement, seek prompt realization of the societal and

environmental benefits of these latest technologies. They are leading efforts to ensure that the best policy decisions are being made for plant (and animal) breeding in Europe and worldwide.

David Hume *National Institutes of Bioscience, UK.*
david.hume@roslin.ed.ac.uk

Marine protection is a UK priority

Scotland's government has responsibility for nature conservation in 61% of UK waters. As programme manager for the Marine Conservation Society in Scotland, I agree that the United Kingdom and its devolved governments should be making marine conservation a priority (see also *Nature* **520**, 585–586; 2015).

This summer, Scotland could become a global conservation leader following proposals for Marine Protected Areas (MPAs) for highly mobile species, including minke whales and basking sharks. These areas will soon be subject to public consultation. The hope is that they will then augment the 30 Scottish MPAs designated in summer 2014. Even at some of these sites, damaging fishing activity may still be continuing (see go.nature.com/hzkdzd). Present and future governments in Scotland and Westminster must take steps to keep the 'P' in MPAs.

Calum Duncan *Marine Conservation Society, Edinburgh; and Scottish Environment LINK, Perth, UK.*
calum.duncan@mcsuk.org

CORRECTION

The Correspondence 'Interim initiative for health in Iran' (*Nature* **521**, 32; 2015) wrongly implied that A. A. Velayati is a volunteer for non-profit group IAPA; only H. Jamaati and S. M. Hashemian are IAPA volunteers.

ARTIFICIAL INTELLIGENCE

Robots with instincts

An evolutionary algorithm has been developed that allows robots to adapt to unforeseen change. The robots learn behaviours quickly and instinctively by mining the memory of their past achievements. [SEE LETTER P.503](#)

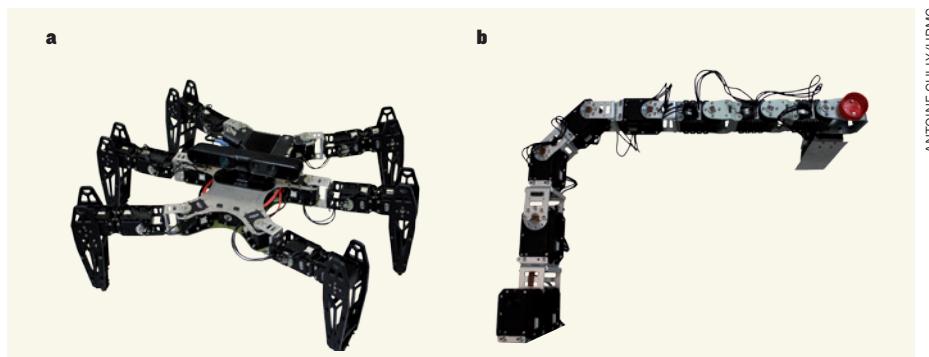
CHRISTOPH ADAMI

Intelligence, by some accounts, is synonymous with the ability to predict the future. Because doing so quickly can often mean the difference between life and death, our brains have evolved to be able to search the vast number of potential futures easily. How is such a feat accomplished? On page 503 of this issue, Cully *et al.*¹ attempt to answer this question by demonstrating that robots can learn to recover quickly and robustly from physical damage — a sudden event that requires them to adopt a new behavioural strategy to continue functioning. The robots (a six-legged mobile robot and a robotic arm; Fig. 1) use a trial-and-error algorithm that lets them tap into the experiences they have accumulated over a simulated lifetime, to quickly find optimal compensating behaviours as if by instinct (see Supplementary Video 1 in ref. 1).

Accurate prediction of events in complex environments requires experience, an understanding of ‘how the world works’, and the capacity to evaluate one’s own actions in the context of those of others. It can be argued that the further out in time an organism or a machine can make accurate predictions of the future, the more intelligent it is. Using this definition, even simple organisms have some intelligence: microbes such as *Escherichia coli*, for example, make predictions about where they must move to find higher concentrations of sugars, and squirrels anticipate the winter by stashing away nuts.

Among animals, humans have arguably the highest level of intelligence, because we can anticipate events hundreds, thousands or even millions of years in the future — albeit largely in domains that do not involve the actions of people, such as planetary orbital dynamics. How can we begin to understand the cognitive underpinnings of a predictive capability that is, to a smaller or larger extent, inherent in all forms of life on Earth? One way, following Richard Feynman’s dictum “What I cannot create I do not understand”, is to recreate intelligence in a machine or robot.

Attempts to create robot intelligence have come and gone with limited success in the past half-century, and it seems as if the goal of creating a machine with human-like intelligence



ANTOINE CULLY/UPMC

Figure 1 | Adaptive machines. Cully *et al.*¹ have designed an algorithm that allows robots to develop strategies for overcoming the effects of damaged limbs. Two robots were used: **a**, a hexapod (width 50 centimetres); **b**, a robotic arm (length 62 cm).

remains elusive — even as great strides are being made. Notwithstanding the successes of chess-playing programs, IBM’s artificially intelligent computer Watson, and the advent of algorithms for self-driving cars, true robot intelligence still eludes us.

Previous studies in the field of robot cognition^{2,3} have suggested that the ability to plan future actions hinges on the ability to recreate a model of the world inside the robot brain — an abstract version, but one that is accurate enough for mental trials and errors to quickly reveal the best strategy to adopt. But even supposing that these model representations⁴ can be generated, how can the vast ‘space’ of likely future actions be searched quickly and efficiently?

Cully *et al.* subjected their robots to several different unforeseen changes in the machines’ morphology (akin to damage), and then asked them to find movement strategies that would compensate for the injury. Before being injured, the robots used an algorithm to establish a baseline of possible actions, which they used after injury to try out moves that were likely to be successful before deciding on any particular compensatory behaviour. Even though the range of possible behaviours (the behaviour space) for a robot might theoretically be infinite, this baseline can be established because, in reality, a robot’s actions are constrained by its morphology.

A hexapod robot such as that studied by the authors is controlled by 36 parameters, but most of the strategies (sequences of motor

activations in a 36-dimensional space) make no sense. Within the robot’s ‘embodiment’⁵ — the way in which the robot’s body is realized — only a small subset of activations can follow any particular prior activation. In other words, the robot’s embodiment dramatically reduces the number of potential strategies, so that sensible actions occupy a severely reduced behaviour space (think of a line instead of a sphere). This reduced space is actually searchable in real time.

The authors created the set of all possible behaviours by having each robot perform many thousands of motions (sequences of motor activations) and recording the ‘fitness values’ of each sequence. The fitness could be as simple as the distance travelled by the robot. Collating this database is time-intensive, but it is analogous to what happens in the natural world, in which living organisms have a lifetime to acquire such information. The robots synthesized new behaviours from this data set using a set of special-purpose machine-learning algorithms that assume that — even in changed circumstances — the actions that are most likely to succeed are ‘close’ to other such actions in a suitably defined behaviour space.

Although these machine-learning algorithms are unlikely to be similar to those used, for example, by mammalian cognitive systems, they share a common premise: that a behaviour space that is dramatically reduced through embodiment, and that is learned from experience, can be searched quickly through

trial and error. If we return to the analogy of a one-dimensional line as opposed to a sphere of possible strategies, only two directions have to be attempted for the line before the preferred direction is clear, whereas in a sphere six directions must be sampled. Given that the robot's behaviour space is 36-dimensional, it is clear that the 'flattening' of the space of options can have dramatic effects.

Could these intuitive trial-and-error strategies be used to discover more-general problem-solving methodologies, of the kind that require planning in uncertain environments? It is difficult to imagine that the method could easily be scaled up to such a level; this particular algorithm was hand-designed by the authors, whereas the 'algorithm' our brains use is the result of millions

of years of Darwinian tinkering and pruning.

Given the failure of past efforts to design robots that display the quick, intuitive and situation-appropriate behaviour of even the smallest rodents, perhaps it is time to give up on the idea that we can design brains, and instead place our hopes in the power of adaptive and evolutionary algorithms. Indeed, the core algorithm that generates the map of possible high-performance behaviours in Cully and colleagues' study is inherently evolutionary, because good strategies are improved on by replication with variation, and selection.

We may never understand our brains in terms of information-processing concepts, but we do understand how to harness the power of evolution. We should therefore let evolution

create for us what we do not understand, one more time. ■

Christoph Adami is in the Department of Microbiology and Molecular Genetics, and the Department of Physics and Astronomy, Michigan State University, Michigan 48824, USA.

e-mail: adami@msu.edu

- Cully, A., Clune, J., Tarapore, D. & Mouret, J.-B. *Nature* **521**, 503–507 (2015).
- Bongard, J., Zykov, V. & Lipson, H. *Science* **314**, 1118–1121 (2006).
- Adami, C. *Science* **314**, 1093–1094 (2006).
- Marsteller, L., Hintze, A. & Adami, C. *Neur. Comput.* **25**, 2079–2107 (2013).
- Pfeifer, R. & Bongard, J. *How the Body Shapes the Way We Think: A New View of Intelligence* (MIT Press, 2007).

CELL BIOLOGY

Polarized transport in the Golgi apparatus

Proteins can be transported in either direction across a cellular organelle called the Golgi apparatus. It emerges that CDC42, a molecule that confers cell polarity, acts to control the directionality of transport in the Golgi. **SEE LETTER P.529**

AKIHIKO NAKANO

Polarity is a part of life on all scales, from head-to-tail polarity in whole organisms to polar orientation in cells. Even intracellular organelles, such as the Golgi apparatus, can be polarized. Proteins synthesized by another organelle, the endoplasmic reticulum, mainly enter the Golgi on one side — dubbed the *cis* side — and, after appropriate modifications, exit from the opposite, *trans*, side. On page 529 of this issue, Park *et al.*¹ identify CDC42 as a protein that regulates the directionality of polarized protein transport within the Golgi apparatus.

The Golgi apparatus, which prepares proteins for shuttling to their final destination in the cell, is a mysterious organelle. More than 100 years have passed since its discovery, but there are still many controversies regarding the mechanism that underpins its function. One debate famous among cell biologists concerns how cargoes such as secretory proteins are transported across the polarized structure of the Golgi from *cis* to *trans*² — the anterograde direction. Various models have been put forward to explain this directionality, but several lines of evidence^{3,4} support 'cisternal maturation'. In this model, flattened, membrane-bound sacs called cisternae, which make up the Golgi, form on the *cis* side of the organelle and move towards the *trans* side as they mature,

carrying proteins produced by the endoplasmic reticulum with them as they go.

However, this mechanism alone cannot

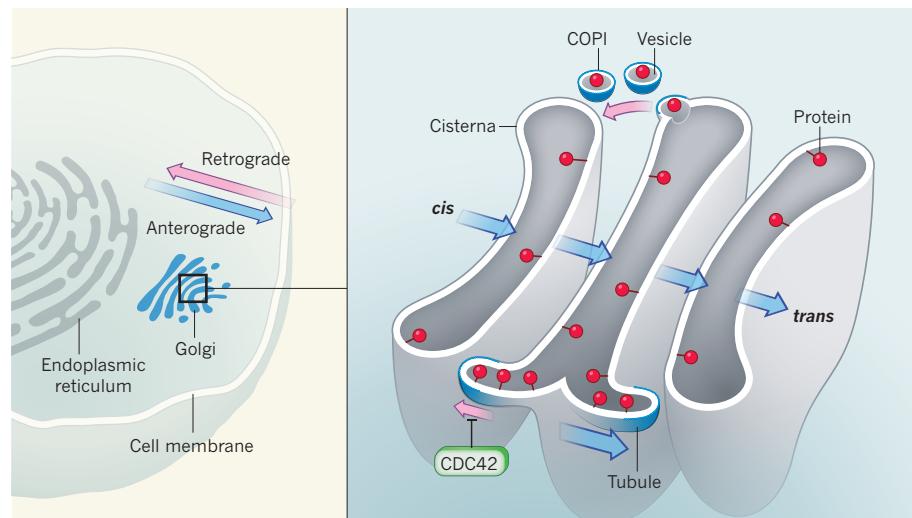


Figure 1 | Bidirectional transport. The Golgi apparatus, which contains a stack of membrane-bound sacs called cisternae, prepares proteins for shuttling around the cell. Proteins mainly move through the Golgi in the anterograde direction, from the *cis* to the *trans* side, travelling from their site of synthesis in the endoplasmic reticulum towards the cell membrane. However, proteins can also move in the opposite, retrograde, direction. Proteins can be transported across the Golgi in three ways: through 'cisternal maturation', in which protein-containing cisternae move from *cis* to *trans* as they mature; in vesicles coated in coat protein complex I (COPI), which mediate retrograde transport; and in tubules, in which COPI binds to proteins to promote transport in either direction. Park *et al.*¹ report that the CDC42 protein competes with COPI for binding of retrograde, but not anterograde, cargoes in tubules, and so enhances anterograde and inhibits retrograde transport.

anterograde and retrograde directions, recognizing different structural signals on each type of cargo. Binding seems to be required for efficient transport of both types of cargo.

Small GTPase proteins act as molecular switches for a wide variety of cellular functions. One small GTPase, CDC42, was originally documented to control cell division in budding yeast, and has since been found to regulate cell polarity^{6,7}. The authors investigated the possibility that CDC42 regulates the selectivity of COPI binding, because it is already known to interact with COPI to modulate anterograde cargo transport in the Golgi⁸. They found that an activated mutant form of CDC42 accelerated anterograde transport in the Golgi, but repressed retrograde transport. By contrast, reducing levels of CDC42 activity inhibited anterograde, but enhanced retrograde, transport.

Next, the authors analysed the role of CDC42 in transport through the Golgi tubules specifically. Using a ‘reconstitution system’ designed to generate COPI-coated tubules⁵, Park *et al.* showed that the activated form of CDC42 reduced the level of retrograde transport by preventing COPI from binding to retrograde cargoes in tubules. However, the mutant protein had no effect on anterograde cargoes. The researchers demonstrated that CDC42 imparts sorting specificity to Golgi tubules in two ways: first, through two lysine amino-acid residues in its carboxy-terminal end that compete with retrograde, but not anterograde, cargoes for binding to COPI; and second, by promoting the formation of tubules at the expense of COPI-coated vesicles, thus reducing retrograde transport. Finally, an analysis of small GTPases, including RhoA, Rac1 and some Rab and Arf GTPases, revealed no others that could affect the directionality of intra-Golgi transport, demonstrating the specificity of CDC42.

If, as these results indicate, CDC42 does indeed modulate the sorting of anterograde and retrograde cargoes in the Golgi tubules, then this study takes an important step towards providing a fuller understanding of the exact mechanism that underpins intra-Golgi transport. CDC42 regulates cellular polarity through its intimate interaction with actin filaments, which help to maintain the structure of the cell⁷. The role of CDC42 in the Golgi tubules probably does not directly involve actin, and so how the protein regulates directionality in this setting is a question for the future. Other, as-yet-undefined factors must be acting with CDC42.

Whether the tubules containing anterograde cargoes do indeed connect Golgi cisternae in anterograde transport also remains to be definitively proved. The authors suggest that SRC tyrosine kinase, an enzyme frequently mutated in cancer, modulates CDC42, which leads them to hypothesize that there might be a relationship between intra-Golgi transport

and tumour development. But, ultimately, the cue that switches on anterograde transport remains elusive.

CDC42 is evolutionarily conserved, being found in organisms from yeast to mammals. Plants also have many Rho-family GTPases. Whether these related proteins play similar parts to CDC42 in cargo binding, tubule formation and directionality in the Golgi would be of great interest to those working on the organelle. Trafficking pathways are thought to have undergone parallel evolution — have these proteins remained fundamental players? ■

Akihiko Nakano is in the Department of Biological Sciences, Graduate School of Science, The University of Tokyo, Tokyo

OCEAN SCIENCE

The origins of a climate oscillation

An index of water-circulation strength in the North Atlantic Ocean has been derived from sea-level measurements. This provides fresh evidence of the ocean’s leading role in multidecadal climate variability. SEE LETTER P.508

SERGEY K. GULEV & MOJIB LATIF

The Earth has warmed considerably during the twentieth and twenty-first centuries, most probably because of the effects of greenhouse gases emitted as a result of human activities. But strong natural variability on a wide range of timescales, from monthly to multidecadal, is superimposed on the long-term global-warming trend and also causes considerable spatial variation of that trend. We therefore need to understand this variability to discriminate anthropogenic effects from natural climate forcing. On page 508 of this issue, McCarthy *et al.*¹ report a method that enables the effects of ocean circulation on one of the most prominent examples of long-term climate variability — the Atlantic Multidecadal Oscillation (AMO) — to be identified from long-term sea-level data.

Natural climate variability can be generated internally by interactions within or between climate-system components such as the atmosphere, ocean and sea ice, and externally by factors such as volcanic eruptions. The AMO is an example of long-term climate variability associated with the ocean². It represents quasi-periodic oscillations of sea-surface temperature (SST) in the North Atlantic Ocean that have a period of about 70 years. The AMO has been described from instrumental SST

113–0033, Japan, and in the Live Cell Super-Resolution Imaging Research Team, RIKEN Center for Advanced Photonics, Wako. e-mail: nakano@riken.jp

- Park, S.-Y., Yang, J.-S., Schmid, A. B., Soberman, R. J. & Hsu, V. W. *Nature* **521**, 529–531 (2015).
- Pelham, H. R. B. & Rothman, J. E. *Cell* **102**, 713–719 (2000).
- Glick, B. S. & Nakano, A. *Annu. Rev. Cell Dev. Biol.* **25**, 113–132 (2009).
- Nakano, A. & Luini, A. *Curr. Opin. Cell Biol.* **22**, 471–478 (2010).
- Yang, J.-S. *et al.* *Nature Cell Biol.* **13**, 996–1003 (2011).
- Johnson, D. I. & Pringle, J. R. *J. Cell Biol.* **111**, 143–152 (1990).
- Hall, A. *Biochem. Soc. Trans.* **33**, 891–895 (2005).
- Wu, W. J., Erickson, J. W., Lin, R. & Cerione, R. A. *Nature* **405**, 800–804 (2000).

This article was published online on 7 May 2015.

records² going back to the mid-nineteenth century and reconstructed from proxy data³ for the past few centuries.

Climate models^{4,5} and analysis of surface heat fluxes⁵ have suggested that the AMO is an internal mode of climate variability originating from changes in the circulation of the Atlantic Ocean, but its origin is still debated. Signals seen in SST may derive from changes in the ocean’s interior that drive circulation changes⁶, or from the influence of factors outside the ocean⁷. Lack of data from the ocean’s subsurface hinders attempts to quantify the relative contributions of internal and external processes to the AMO.

McCarthy and colleagues’ method for identifying how ocean circulation affects the AMO is based on the hypothesis that ocean currents on relatively large scales of space and time (several tens of kilometres and several weeks) are broadly geostrophic — that is, the direction and strength of the currents depends on the balance between the Coriolis force associated with Earth’s rotation and the horizontal pressure-gradient force in the ocean (Fig. 1). Their approach builds on the proposal⁸ that the difference in dynamic height (the sea level caused by variations in the ocean’s depth-integrated density) between the centres of large, primarily wind-driven, ocean-circulation systems called gyres in the subtropical and subpolar North Atlantic provides a proxy for long-term

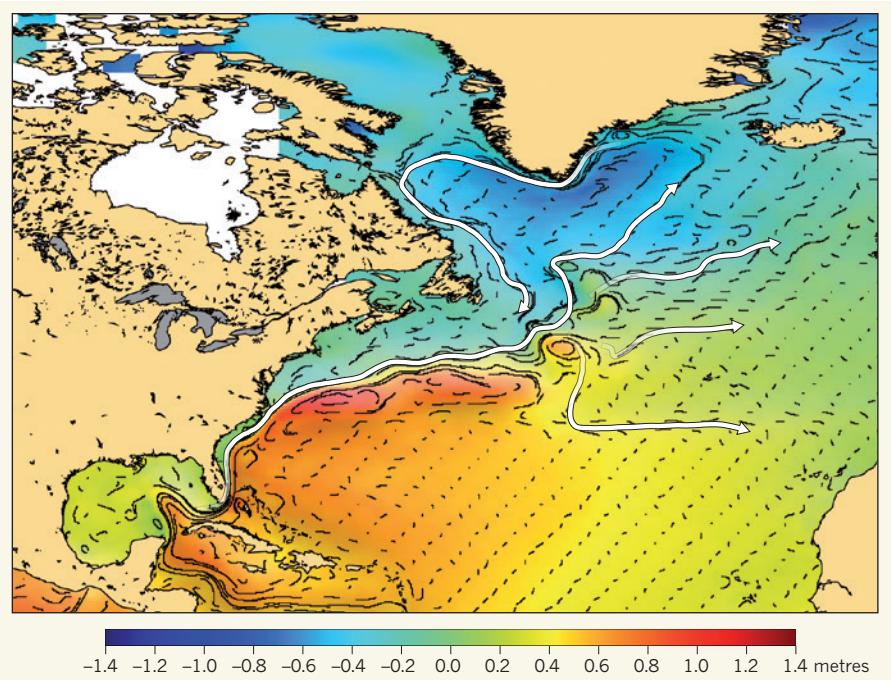


Figure 1 | Circulation and pressures in the North Atlantic Ocean. The map depicts the mean dynamic topography — the oceanic relief that corresponds to long-term, averaged ocean circulation — in the North Atlantic for the years 1993 to 2013. Positive and negative values represent heights above and below the level of the unperturbed ocean, respectively. The topography is a measure of ocean pressure: low pressure occurs in the subpolar-gyre region (blue) and high pressure in the subtropical-gyre region (red). Short black lines indicate local current velocity vectors and white arrows indicate the main ocean currents, which follow maximum pressure gradients. McCarthy *et al.*¹ report a method for identifying how ocean circulation affects the Atlantic Multidecadal Oscillation — quasi-periodic oscillations of sea surface temperature in the North Atlantic that have a period of about 70 years. Their approach assumes that the difference in heights between the centres of the gyres provides a proxy for long-term changes in mass transport in the North Atlantic. Figure adapted from CNES/CLS 2012 (produced by AVISO).

changes in mass transport by the Gulf Stream and its extension, the North Atlantic Current. The difference in heights is a measure of the pressure gradient between the gyres' centres and has previously been estimated⁸ from offshore sea-level records in the subtropics at Bermuda, and from long-term hydrographic data at subpolar latitudes.

But the Bermuda sea-level data are noisy; coastal tide gauges would provide more-robust signals^{1,8}. By taking sea-level differences between coastal tide gauges to the south and north of Cape Hatteras — the boundary between the subtropical and subpolar gyres — as an estimate of the pressure gradient between the gyres, McCarthy and co-workers derived a simple index of the North Atlantic circulation strength (see Fig. 1 of the paper¹) that closely correlates with heat transport in ocean models. Positive values of this index reflect an enhanced pressure gradient compared with the long-term mean and imply strong northward heat transport and increasing upper-ocean heat content in mid-latitude and subpolar regions of the North Atlantic. Negative values imply lowering of the pressure gradient, weakening of northward heat transport and diminishing heat content in the upper ocean. Although this index is

influenced by many other factors in the Gulf Stream region, it seems to be an effective proxy for the AMO.

McCarthy *et al.* also find that fluctuations in the sea-level index (and therefore changes in ocean circulation) precede anomalies in upper-ocean heat content in the subpolar-gyre region by approximately two years. This time lag is important, because it potentially enables prediction of heat-content changes in the subpolar-gyre region from sea-level data, and thus prediction of AMO evolution and related climate anomalies.

The authors also used a numerical model of the ocean (an eddy-permitting ocean general circulation model, for those in the know), forced by observed atmospheric conditions at the surface, to investigate the interconnections between dynamic sea-level gradients, circulation strength and upper-ocean heat content, including time lags. Despite the limitations of such models, especially in the Gulf Stream region and at high latitudes, the simulations largely support the results from the data (Fig. 2 of ref. 1), which is reassuring.

McCarthy and colleagues' study supports previous conjectures⁶ that the ocean integrates chaotic atmospheric variability primarily associated with a climatic phenomenon known as

the North Atlantic Oscillation⁹, and responds by generating SST variability on multidecadal time scales through dynamical ocean processes and with a time lag of about a decade — specifically by causing changes in the large-scale ocean circulation called the Atlantic Meridional Overturning Circulation (AMOC). The findings therefore further reinforce the idea that the AMO is an internal mode of climate variability.

The work also has crucial implications for our understanding of the ocean's role in climate variability and change. First, if the AMO originates from ocean dynamics, then attempts to detect early signs of the effects of anthropogenic climate change in the North Atlantic will be hampered. For example, the climate effects of AMOC slowdown projected by many climate models in response to global warming could be masked for decades by a positive phase of the AMO. Second, the AMO has been linked to climate anomalies in many regions of the globe^{10,11}. The study therefore suggests that decadal predictions of climate over the North Atlantic and adjacent continents are possible if enough ocean data are available for forecast initialization. This would be of enormous societal benefit to many countries.

Further efforts are needed to conclusively prove that the ocean steers the AMO. Relevant instrumental records (including sea levels from tide gauges) are available only for the past several decades, with the longest dating back to the 1920s and 1930s. But a thorough assessment of the ocean's role in the multidecadal climate variability of the North Atlantic sector will require proxy data dating back for many centuries or even millennia. Multi-year experiments to measure ocean currents and cross-sections — using moored measuring devices, satellite data, and buoys that sample the upper ocean — have enabled studies over the past two decades and have captured long-term data. These programmes must continue, given the long timescales involved in the AMO. Climate models initialized with such data will provide more-reliable predictions of future climate than are currently available.

The synthesis of models and data will potentially allow the prediction of multidecadal variations. However, climate models still suffer from large biases, especially in the North Atlantic. Improvements to climate models are therefore also urgently needed. But it remains to be seen whether atmosphere–ocean interactions such as those described by McCarthy *et al.* will remain the same under anthropogenic climate change. ■

Sergey K. Gulev is at the P. P. Shirshov Institute of Oceanology, Russian Academy of Sciences, Moscow 117997, Russia, and at Moscow State University. **Mojib Latif** is at the GEOMAR Helmholtz-Zentrum für Ozeanforschung Kiel, Kiel D-24105,

Germany, and at Kiel University.
e-mail: gul@sail.msk.ru

1. McCarthy G. D., Haigh, I. D., Hirschi, J. J.-M., Grist, J. P. & Smeed, D. A. *Nature* **521**, 508–510 (2015).
2. Knight, J. R., Allan, R. J., Folland, C. K., Vellinga, M. & Mann, M. E. *et al.* *Geophys. Res. Lett.* **32**, L20708 (2005).

3. Delworth, T. L. & Mann, M. E. *Clim. Dynamics* **16**, 661–676 (2000).
4. Latif, M. *et al.* *J. Climate* **17**, 1605–1614 (2004).
5. Gulev, S.K., Latif, M., Keenlyside, N., Park, W. & Koltermann, K. P. *Nature* **499**, 464–467 (2013).
6. Latif, M. in *Ocean Circulation and Climate* 2nd edn (eds Siedler, G., Griffies, S., Gould, J. & Church, J.) (Academic, 2013).
7. Booth, B. B. B., Dunstone, N. J., Halloran, P. R.,

- Andrews, T. & Bellouin, N. *Nature* **484**, 228–232 (2012).
8. Curry, R. G. & McCartney, M. S. *J. Phys. Oceanogr.* **31**, 3374–3400 (2001).
9. Hurrell, J. W. *Science* **269**, 676–679 (1995).
10. Sutton, R. T. & Hodson, D. L. R. *Science* **309**, 115–118 (2005).
11. Zhang, R. & Delworth, T. L. *Geophys. Res. Lett.* **33**, L17712 (2006).

CANCER METABOLISM

A waste of insulin interference

Many people with cancer die from a wasting disorder called cancer-associated cachexia. Two studies in fruit flies show that inhibition of insulin signalling causes cachexia-like organ wasting.

ERWIN F. WAGNER & MICHELE PETRUZZELLI

It is estimated that up to 30% of people with advanced-stage cancer are killed not by the tumour itself, but by a metabolic disorder called cancer-associated cachexia (CAC)^{1,2}, which is characterized by systemic inflammation, weight loss, body-fat atrophy and skeletal-muscle wasting. Cachexia is also a feature of several chronic conditions^{2,3}, including heart failure, lung disease and infectious diseases such as HIV. At present there is no cure for CAC and no biomarkers are available for

identifying patients at risk of developing the disorder. Thus, there is a need to better understand the origins of cachexia, its systemic progression and the molecular pathways involved in its development. Two papers^{4,5} in *Developmental Cell* now report that interference with insulin signalling in the fruit fly *Drosophila melanogaster* induces systemic organ wasting reminiscent of human cachexia.

The cancer-causing protein Yap1 is part of the Hippo signalling pathway, which induces cell proliferation, in part by increasing signalling by insulin and by a related protein,

insulin-like growth factor 1 (IGF1) (ref. 6). Kwon *et al.*⁴ investigated the systemic changes caused by intestinal activation of Yorkie, an equivalent of Yap1, in fruit flies. The authors report that activation of Yorkie in the gut causes over-proliferation of intestinal cells, and leads to the secretion of ImpL2, an insulin growth factor binding protein (IGFBP) that inhibits insulin and IGF1 signalling. Secretion of ImpL2 in the intestine caused systemic wasting in muscles and distant organs, including the ovaries and the fat body (an organ analogous to vertebrate fat and the liver). This remarkable finding implies that over-proliferation in a given tissue can lead to distant metabolic changes and wasting symptoms, similar to human cachexia.

Using a different fly model, Figueroa-Clarevega and Bilder⁵ find evidence to corroborate Kwon and colleagues' study. They transplanted tumours into adult flies, and found that malignant, but not benign, tumours caused wasting of fat, muscle and gonadal tissue. Furthermore, the tumours induced down-regulation of the insulin signalling pathway in peripheral organs, leading to insulin resistance in these regions. Investigating the differences between benign and malignant tumours, the authors found that *ImpL2* was one of the genes most upregulated in malignant tumours. They demonstrated that overexpression of *ImpL2* in specific tissues in flies without tumours led to wasting in distant locations. However, inhibition of *ImpL2* in flies with tumour transplants only partially ameliorated wasting, suggesting that other molecular interactions between the tumour and the host remain to be discovered.

Notably, the systemic wasting traits described in the two papers correlate with other features that are associated with high blood sugar. Furthermore, this wasting is independent of changes in food intake or local damage to organs at the site of tumour growth. This, taken together with the other results, indicates that factors secreted from tumours or hyper-proliferating tissues in *D. melanogaster* perturb systemic metabolism and induce wasting in distant organs.

Insulin resistance is common in patients with cancer, and may contribute to skeletal-muscle wasting in mouse models of CAC⁷. Previous studies⁸ of how tumour-derived factors affect the organism have primarily focused on a group of cell-signalling molecules called pro-inflammatory cytokines, which promote systemic inflammation. Alterations in systemic

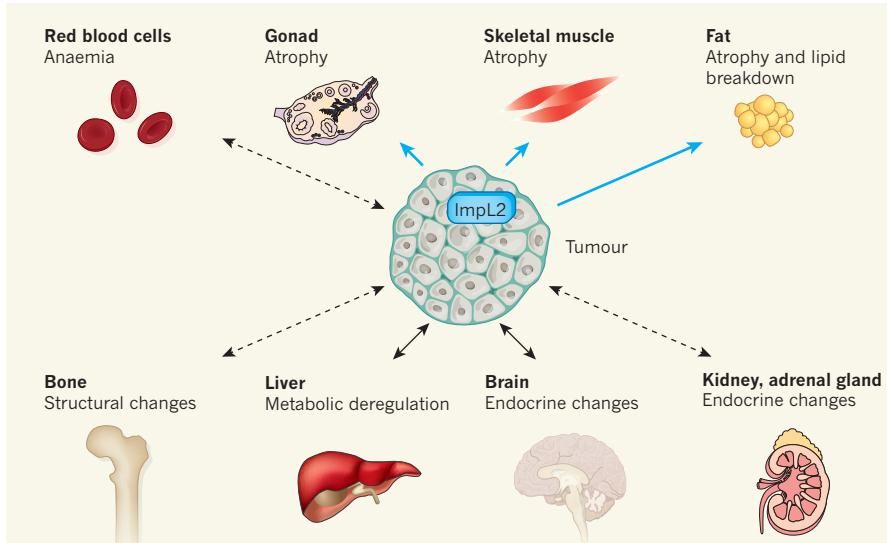


Figure 1 | Systemic metabolic changes in cancer. Cancer-associated cachexia is defined by harmful abnormalities that are brought about by factors secreted by the growing tumour, and by adaptive responses in host tissues. Solid double-headed arrows indicate that the molecular mechanisms underpinning the trait (which are thought to be bidirectional) have been defined, whereas dashed double-headed arrows indicate that causative factors remain unknown. Kwon *et al.*⁴ and Figueroa-Clarevega and Bilder⁵ demonstrated that the protein ImpL2, which is secreted from hyperproliferating tissues and tumours, impairs insulin signalling in fruit flies. This interference is the cause of atrophy in skeletal muscle, fat and gonadal tissues (indicated by the blue arrows, which are unidirectional to reflect the fact that bidirectional interactions have little time to become established in the fly studies.)

metabolism that arise as a consequence of tumour growth have been regarded as secondary effects of inflammation. The current studies provide evidence that impaired insulin signalling is itself a direct cause of CAC development. However, whether insulin resistance is sufficient to cause cachexia in model systems other than fruit flies remains to be demonstrated.

Insulin and IGF1 signalling are key regulators of tissue mass in both flies and mammals, and it is possible that IGFBPs are differentially regulated in cancers that are associated with cachexia compared with those that are not. Little is known about the regulation of ImpL2 in flies and mice, but it is probable that stress factors, which activate inflammatory pathways and Hippo signalling⁶, could induce ImpL2 expression. It is interesting that other tumour-specific proteins, such as the *Drosophila* cytokine Upd2, are unable to induce organ wasting in flies, whereas the equivalent protein in mice, IL-6, is a mediator of CAC⁸. Surprisingly, a role for the immune system is not

discussed in the two papers despite inflammation being an accepted hallmark of cachexia.

These two studies highlight the importance of studying the metabolic response to cancer. Although our knowledge of the metabolism of cancer cells themselves is steadily improving, the characterization of organism-wide metabolic changes in response to cancer is still incomplete (Fig. 1). The focus of cancer research is gradually expanding, from the cancer cell to the tumour microenvironment, to the system as a whole. Supporting the need to study organism-wide metabolism, abnormal alterations in organs at a distance from the primary tumour that are independent of the process of cancer metastasis have been described^{8–10}. Remarkably, targeting such alterations has therapeutic value in mice, ameliorating total body-weight loss and skeletal-muscle and adipose-tissue atrophy, without directly affecting the mass of the tumour^{8–10}. The tiny fruit fly nicely illustrates the value of broadening our horizons to encompass the organism as a whole, and

of using animal models of cancer to explore this macroenvironment. ■

Erwin F. Wagner is in the Cancer Cell Biology Programme, CNIO, Madrid 28029, Spain.

Michele Petruzzelli is in the Department of Oncology, University of Cambridge, Addenbrooke's Hospital, Cambridge CB2 0QQ, UK.

e-mails: ewagner@cnio.es;
mp753@cam.ac.uk

1. Tisdale, M. J. *J. Physiol. Rev.* **89**, 381–410 (2009).
2. Fearon, K., Arends, J. & Baracos, V. *Nature Rev. Clin. Oncol.* **10**, 90–99 (2013).
3. Argilés, J. M., Busquets, S., Stemmler, B. & López-Soriano, F. J. *Nature Rev. Cancer* **14**, 754–762 (2014).
4. Kwon, Y. et al. *Dev. Cell* **33**, 36–46 (2015).
5. Figueiroa-Clarevega, A. & Bilder, D. *Dev. Cell* **33**, 47–55 (2015).
6. Strassburger, K., Tiebe, M., Pinna, F., Breuhahn, K. & Teleman, A. A. *Dev. Biol.* **367**, 187–196 (2012).
7. Fearon, K., Glass, D. J. & Guttridge, D. C. *Cell Metab.* **16**, 153–166 (2012).
8. Petruzzelli, M. et al. *Cell Metab.* **20**, 433–447 (2014).
9. Kir, S. et al. *Nature* **513**, 100–104 (2014).
10. Das, S. K. et al. *Science* **333**, 233–238 (2011).

MICROBIOLOGY

Taking the bad with the good

Modelling of the interactions between antibiotic production and antibiotic degradation reveals that these opposing activities are key to maintaining diversity in microbial communities. SEE LETTER P.516

CARL T. BERGSTROM & BENJAMIN KERR

We commonly expect competitive ecological interactions to be transitive: if ravens displace crows and crows displace jays, then ravens should displace jays as well. But the world does not always work this way. Increasingly, researchers are finding intransitive relationships, in which ravens displace crows, which displace jays, which in turn displace ravens. Intransitive relationships occur in animals^{1,2}, plants^{3,4} and microbes^{5,6}. Theoretical models show that species abundances can cycle in intransitive communities, in principle preserving species diversity^{7,8}. However, in finite populations, extinction can readily occur when one type cycles to low abundance. On page 516 of this issue, Kelsic *et al.*⁹ model an intransitive system in which microbial species produce antimicrobial compounds and exhibit differing sensitivities to the products of their competitors. By demonstrating that antimicrobial degradation can stabilize a multi-species community, the authors suggest a new solution to the puzzle of how bacterial diversity is maintained¹⁰.

To illustrate Kelsic and colleagues' model, consider the rock–paper–scissors (RPS) scenario familiar to many as a game. Imagine three microbial species called Rock, Paper and Scissors, each of which produces a unique antimicrobial compound and is immune to its own toxin. If Rock kills Scissors, Scissors kills Paper and Paper kills Rock, we have the standard situation (Fig. 1a). Each bacterial species must protect itself from the toxin of its victim; for example, Scissors protects itself from Paper's toxin. Kelsic *et al.* focused on a neglected aspect of this protection: it may be non-excludable, meaning that protection may spill over to other species.

Such ‘leakiness’ may occur if a cell degrades the antimicrobials of a competing species by secreting enzymes that do the job externally, or by deactivating the competitor's antimicrobials once they have entered the cell¹¹. Either way, the concentration of the antimicrobial in the environment is reduced. The RPS scenario can be adapted to account for this leakiness. For instance, when Scissors protects itself from Paper's toxin, partial protection would extend to Rock as well — here, Scissors inadvertently

helps its own enemy (Fig. 1b, orange line). Kelsic and colleagues develop a mathematical model showing that the ‘public good’ of leaky protection and the ‘public bad’ of toxin production can interact to permit stable coexistence between multiple species.

Models with diffusible public goods and public bads are complicated. To improve our intuition about Kelsic and colleagues' model, let us reframe this population-level interaction as a two-player game. Instead of having many toxic microbes interacting with one another simultaneously, we consider pairwise interactions between individuals, each of whom play one of the RPS strategies. Players meet at random and receive a pay-off of 1 for winning or drawing and 0 for losing. They then replicate according to their pay-offs, and faithfully pass on their strategies to their offspring. A population of individuals playing this game undergoes unstable cycles; this scenario corresponds to the mathematical model developed by Kelsic *et al.* when there is no leaky protection. If instead the pay-off for winning is 2, drawing is 1 and losing is 0, we have the RPS game more commonly studied in the literature; this version has neutrally stable cycles⁸. Whether the cycles are neutrally stable or unstable, two of the three strategies will eventually be lost in a finite population (Fig. 1a).

Now suppose that we add one or more ‘bystanders’ to the game. Within the pair, a would-be winner is ineffectual if its enemy is standing by, and the game ends in a draw. With one or two bystanders, randomly chosen from the population, the chances of having a bystander that can interfere are too low to alter the dynamics qualitatively. But with more bystanders, the dynamics change completely: the community can approach a stable balance

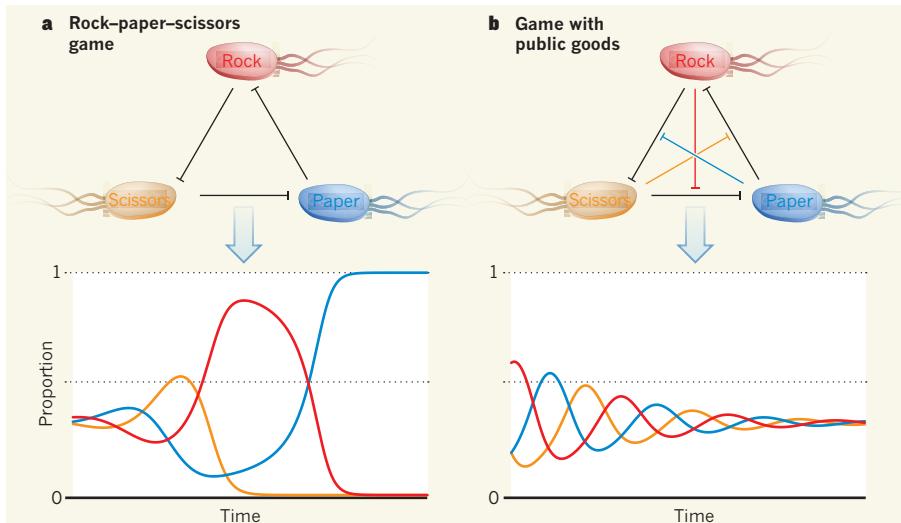


Figure 1 | Cyclical dynamics in a ‘rock–paper–scissors’ game with public goods. **a**, In a standard rock–paper–scissors game, rock beats scissors, which beats paper, which beats rock. Here, Rock, Paper and Scissors represent bacterial species, each of which produces an antimicrobial toxin (black lines). Each species is immune to the toxin produced by its victim. Kelsic *et al.*⁹ show that, when such immunity is not ‘leaky’, the proportions of the three species fluctuate in unstable cycles until two of the species are lost. **b**, When immunity is conferred by antibiotic degradation, this protection may spill over to other species (coloured lines). When such ‘public goods’ are produced in appropriate quantities, the cycles dampen in amplitude and all three types coexist stably.

of the three strategies, in which all three types can coexist indefinitely (Fig. 1b). The more bystanders that are present, the easier it is for coexistence to be maintained.

Why do bystanders facilitate coexistence? In the standard RPS game, each type directly hurts its victim and thereby indirectly helps its enemy (its victim’s victim). This process occurs in the bystander version of the game as well, but in addition, each type directly helps its enemy. As a player, Paper defeats a Rock partner; and as a bystander, Paper prevents Rock from defeating Scissors. Thus, Paper helps Scissors in two ways: by reducing the proportion of Rock players overall (the indirect route) and by ameliorating the harm that Rock does to Scissors when they meet (the direct route). If one type helps its enemy enough by the direct route, the additional source of feedback halts the proliferation of any type that becomes common, and can stabilize the entire community.

By adding bystanders to the basic RPS game, the model captures the essential biological feature of bacterial interactions that drives Kelsic and colleagues’ findings: that bacteria of a first species protect bacteria of a second one when they protect themselves by disabling the antimicrobial compound of a third. Thus, addition of bystanders captures the leakiness of antimicrobial protection. (For a full description of the bystander model and its dynamics, see go.nature.com/yjzkfm.)

Kelsic *et al.* make several important contributions to our understanding of microbial community dynamics. They demonstrate experimentally that one microbe can protect a second from a third. They demonstrate mathematically how bystander protection can

stabilize multi-species bacterial communities. They show that their qualitative findings are robust to variation in the form of the model, its parameters, the community structure and the number of species. Finally, they illustrate that the evolution of new strategies need not disrupt stable communities of this sort. In particular, coexistence can be robust against ‘cheaters’: mutants that obtain a fitness benefit by ceasing to produce public goods or public bads.

The authors also acknowledge several caveats to their work. If bystander protection is excessively weak or excessively strong, species cannot coexist indefinitely. Even if the protection is just right, coexistence will

not occur if the starting proportions of each species are too uneven.

Kelsic *et al.* have derived a key theoretical result: that, in principle, multiple microbial species can coexist indefinitely even in a well-mixed environment without spatial refugia. The obvious next step is to return to experimental studies to test whether this mechanism can support bacterial diversity in practice. Is protection strong enough to stabilize coexistence, and can a well-mixed three-species system coexist stably under laboratory conditions? Coexistence can withstand the emergence of cheaters in the model; can it do so in microbial culture as well? We are eager to learn whether, in natural microbial communities with intransitive competition, diversity is promoted when public goods take out public bads. ■

Carl T. Bergstrom and Benjamin Kerr are in the Department of Biology, University of Washington, Seattle, Washington 98195-1800, USA.

e-mails: cbergst@uw.edu; kerrb@uw.edu

1. Sinervo, B. & Lively, C. M. *Nature* **380**, 240–243 (1996).
2. Jackson, J. B. C. & Buss, L. *Proc. Natl Acad. Sci. USA* **72**, 5160–5163 (1975).
3. Lankau, R. A. & Strauss, S. Y. *Science* **317**, 1561–1563 (2007).
4. Cameron, D. D., White, A. & Antonovics, J. J. *Ecol.* **97**, 1311–1319 (2009).
5. Paquin, C. E. & Adams, J. *Nature* **306**, 368–371 (1983).
6. Kerr, B., Riley, M. A., Feldman, M. W. & Bohannan, B. J. M. *Nature* **418**, 171–174 (2002).
7. Gilpin, M. E. *Am. Nat.* **109**, 51–60 (1975).
8. Frean, M. & Abraham, E. R. *Proc. R. Soc. B* **268**, 1323–1327 (2001).
9. Kelsic, E. D., Zhao, J., Vetsigian, K. & Kishony, R. *Nature* **521**, 516–519 (2015).
10. Kassen, R. & Rainey, P. B. *Annu. Rev. Microbiol.* **58**, 207–231 (2004).
11. Yurtsev, E. A., Chao, H. X., Datta, M. S., Artemova, T. & Gore, J. *Mol. Syst. Biol.* **9**, 683 (2013).

This article was published online on 21 May 2015.

PALAEONTHROPOLOGY

The middle Pliocene gets crowded

New hominin fossils discovered in Ethiopia, dated to between 3.5 million and 3.3 million years ago, suggest that species diversity may have been as high during early human evolution as in later periods. SEE ARTICLE P.483

FRED SPOOR

For many years, human evolution was viewed as a diverse radiation of species emerging a little after 3 million years (Myr) ago following an earlier phase characterized by little or no diversity (Fig. 1). Best

known from this earlier period is the middle Pliocene species *Australopithecus afarensis*, which is documented by fossils between 3.7 and 3.0 Myr old from eastern Africa, including the emblematic skeleton known as Lucy¹. On page 483 of this issue, Haile-Selassie *et al.*² describe a new species, *Australopithecus*

deyiremeda, which adds to other evidence^{3–8} indicating that *A. afarensis* was not the only ancient human relative during this time period. Indeed, it seems that the hominins that populated Africa in the middle Pliocene may have been just as diverse taxonomically as later stages of human evolution.

In the 1970s and 1980s, it was hotly debated whether the middle Pliocene hominin fossils found in Ethiopia and Tanzania represent more than one species, but a broad consensus emerged that they could all be attributed to the single variable species *A. afarensis*¹. Further fossil finds rekindled the debate about taxonomic diversity and led to the naming of two new species: *Australopithecus bahrelghazali* in 1996, based on a 3.6-Myr-old partial lower jaw from Chad³, and *Kenyanthropus platyops* in 2001, based on cranial fossils 3.5–3.3 Myr old from Lomekwi, Kenya⁴. The notion of multiple, contemporary middle Pliocene species was highlighted more recently by the discovery at Burtele, Ethiopia, of 3.4-Myr-old foot bones that are too primitive to belong to *A. afarensis*⁷, and by the revised dating of a South African *Australopithecus* skeleton to 3.7 Myr ago⁸ (Fig. 1).

It is against this background that Haile-Selassie and colleagues describe *A. deyiremeda*, which is represented by fossils between 3.5 Myr and 3.3 Myr old discovered at Woranso-Mille in Ethiopia. The location is close to sites of similar age that have produced abundant *A. afarensis* specimens¹, suggesting that the two species lived in close proximity in place as well as in time.

The *A. deyiremeda* fossils include a partial maxilla (upper jaw bone) and two mandibles (lower jaw bones). The maxilla shows that the cheek bones were situated more forward than in *A. afarensis*, and the mandibles have a thick body because the ramus (the part connecting with the cranium) ascends from a more-forward position. The cheek teeth are notably small, except for the lower third molars. Such comparisons can be made with confidence because *A. afarensis* is known from one of the largest samples in the hominin fossil record, providing a good representation of morphological variation¹.

But how does *A. deyiremeda* compare with the other middle Pliocene hominin species? As noted by Haile-Selassie and colleagues, it shares several diagnostic features with *K. platyops*, including forward cheek bones, small first and second molars, and three-rooted upper premolars⁴. However, the two species differ in the front part of the maxilla, which is flat and non-projecting in *K. platyops* but curved and protruding in *A. deyiremeda* and *A. afarensis*. Similarly, the *A. deyiremeda* mandibles share a forward position of the ramus with two possible *K. platyops* specimens from Lomekwi, but they are more like *A. afarensis* in the midplane. In broad terms, *A. deyiremeda* can be seen as sharing newly evolved morphology with *K. platyops* in the cheek areas of the jaws, while showing a more primitive, *A. afarensis*-like

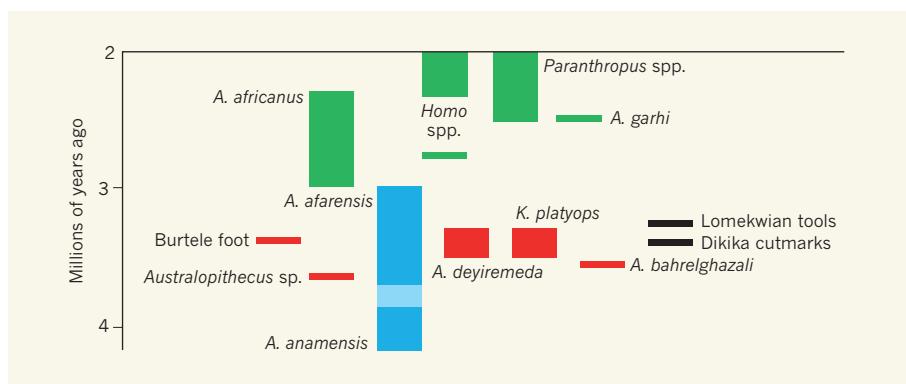


Figure 1 | Temporal distribution of known hominin species. Bars show when evidence of various hominin species occurs in the fossil record between 4 million and 2 million years (Myr) ago. The long-standing interpretation of hominin diversity was that *Australopithecus afarensis* and its probable ancestor *A. anamensis* formed a single lineage during the middle Pliocene (blue bars), and that this was followed by a radiation of *Australopithecus*, *Homo* and *Paranthropus* species after 3 Myr ago (green bars). However, more-recent fossil evidence indicates taxonomic diversity during the middle Pliocene (red bars), including *Kenyanthropus platyops*, multiple species of *Australopithecus*, and an unassigned foot fossil from Burtele, Ethiopia. Haile-Selassie and colleagues' announcement² of the species *Australopithecus deyiremeda* adds to this early diversity. The earliest archaeological evidence for the use of stone tools, from the Lomekwi and Dikika sites, is also indicated.

shape towards the front. Whether the retention of this primitive morphology in *A. deyiremeda* does indeed warrant species-level distinction from *K. platyops* must be confirmed by further study and fossil finds. Comparison with *A. bahrelghazali* is difficult because this species is known only from the front part of a mandible and an isolated premolar. It seems that one or more additional specimens have been recovered^{1,6} but these remain undescribed.

Finding such taxonomic diversity raises the question of how multiple species could have coexisted over a long period in a stable ecosystem, particularly when they live in close geographic proximity, as seems to be the case with *A. deyiremeda* and *A. afarensis*. Niche partitioning, involving diverse dietary preferences, foraging strategies, habitat selection and population movements, will probably be the key factor. However, establishing a concrete link between such characteristics and the morphological differences that distinguish species is often difficult, not least because the morphology may be affected by random genetic drift as much as by selection. *A. deyiremeda*, *A. afarensis* and *K. platyops* differ in their jaws in ways that reflect the position of the chewing muscles and the size of the tooth crowns. These differences provide an opportunity to investigate whether feeding behaviour and diet played a part, by modelling the biomechanics of chewing and assessing the dental wear and stable isotopes present in the fossils, both of which can give an indication of the types of food eaten by the individual.

Any discussion of niche partitioning and foraging behaviour among middle Pliocene hominins should consider the discovery of 3.4-Myr-old cut-marked bones at Dikika⁹ in Ethiopia, and of 3.3-Myr-old stone tools at Lomekwi¹⁰. Differences in the use of stone tools may reflect species-specific behavioural

niches, so it would be informative if we could identify the Lomekwi toolmaker. The tools were found in close geographical and temporal proximity to where a key *K. platyops* specimen (the paratype) was discovered, making this species the default candidate.

However, a further complication now emerges: the *K. platyops* paratype is a maxillary fragment that was associated with the holotype (main defining specimen) on the basis of three shared features: forward cheek bones, three-rooted premolars, and a small first-molar crown⁴. But these three features are now also found in *A. deyiremeda*, and the distinctive front part of the maxilla that sets *K. platyops* apart is not preserved in the paratype. Hence, it cannot be excluded that it was *A. deyiremeda* that made an appearance at the tool site. Regardless, associating stone tools with a specific species as the toolmaker is notoriously difficult, and the increasingly rich fossil record of the middle Pliocene provides plenty of opportunity for lively debate. ■

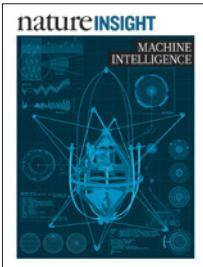
Fred Spoor is in the Department of Cell and Developmental Biology, University College London WC1E 6BT, UK, and the Department of Human Evolution, Max Planck Institute for Evolutionary Anthropology, Leipzig, Germany. e-mail: f.spoor@ucl.ac.uk

1. Kimbel, W. H. & Delezene, L. K. *Yearb. Phys. Anthropol.* **52**, 2–48 (2009).
2. Haile-Selassie, Y. et al. *Nature* **521**, 483–488 (2015).
3. Brunet, M. et al. *C. R. Acad. Sci.* **322**, 907–913 (1996).
4. Leakey, M. G. et al. *Nature* **410**, 433–440 (2001).
5. Spoor, F., Leakey, M. G. & Leakey, L. N. *Phil. Trans. R. Soc. B* **365**, 3377–3388 (2010).
6. Guy, F. et al. *J. Hum. Evol.* **55**, 37–47 (2008).
7. Haile-Selassie, Y. et al. *Nature* **483**, 565–569 (2012).
8. Granger, D. E. et al. *Nature* <http://dx.doi.org/10.1038/nature14268> (2015).
9. McPherron, S. P. et al. *Nature* **466**, 857–860 (2010).
10. Harmand, S. et al. *Nature* **521**, 294–295 (2015).

nature INSIGHT

MACHINE INTELLIGENCE

28 May 2015 / Vol 521 / Issue No 7553



Cover illustration

Nik Spencer

Editor, *Nature*

Philip Campbell

Publishing

Richard Hughes

Production Editor

Jenny Rooke

Art Editor

Nik Spencer

Sponsorship

Reya Silao

Production

Ian Pope

Marketing

Steven Hurst

Editorial Assistants

Rebecca White

Melissa Rose

The Macmillan Building
4 Crinan Street
London N1 9XW, UK
Tel: +44 (0) 20 7833 4000
e: nature@nature.com



nature publishing group

When the phrase ‘artificial intelligence’ was first coined in 1956 by computer scientists, expectations for the development of machines with human-like intelligent reasoning and behaviour were high, but the field was in for a long wait. In the decades to come, the term referred largely to popular culture. Worse, the relative failure of algorithms that attempted to mimic human reasoning at the higher, symbolic level gave the field a bad name and led to a long-term freeze in funding.

In the meantime, conventional, not-so-intelligent computers became faster, more powerful and consumer-friendly thanks to purposeful investments by an ambitious electronics industry.

However, there is not much room left for computer chips to improve further, owing to physical limitations. In the past decade, the world has also seen a massive explosion of data, and retrieving meaningful information from this deluge will soon be impossible with conventional computers. Fortunately, work on artificial neural networks has steadily continued in the background, and the field has recently made big conceptual breakthroughs. Together with the availability of powerful computer processors and large amounts of data for the algorithms to train on, the field of artificial intelligence has made a comeback, demonstrating machine-learning applications such as those that process visual and linguistic information in a human-like manner.

Another route to machine intelligence is robotics, in which the presence of an artificial being in the physical world as well as sensory input is essential to its intelligent behaviour. This area, profiting from advances in chip technology, computer algorithms and smart materials, is making big strides towards the creation of robots that can safely assist humans in a range of tasks.

In this Insight, we have collected some of the most exciting developments in machine learning and robotics. Expectations are again high, but, as the following Reviews demonstrate, there are several exciting avenues now open to further research. With the right safeguards in place, these opportunities could be essential to addressing the challenges of a complex world in the twenty-first century.

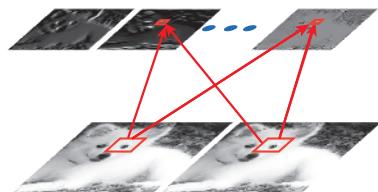
Tanguy Chouard & Liesbeth Venema
Senior Editors

CONTENTS

REVIEWS

436 Deep learning

Yann LeCun, Yoshua Bengio & Geoffrey Hinton



445 Reinforcement learning improves behaviour from evaluative feedback

Michael L. Littman

452 Probabilistic machine learning and artificial intelligence

Zoubin Ghahramani

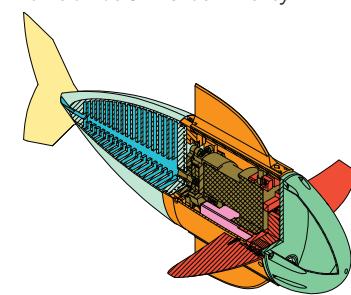
460 Science, technology and the future of small autonomous drones

Dario Floreano & Robert J. Wood



467 Design, fabrication and control of soft robots

Daniela Rus & Michael T. Tolley



476 From evolutionary computation to the evolution of things

Agoston E. Eiben & Jim Smith

Deep learning

Yann LeCun^{1,2}, Yoshua Bengio³ & Geoffrey Hinton^{4,5}

Deep learning allows computational models that are composed of multiple processing layers to learn representations of data with multiple levels of abstraction. These methods have dramatically improved the state-of-the-art in speech recognition, visual object recognition, object detection and many other domains such as drug discovery and genomics. Deep learning discovers intricate structure in large data sets by using the backpropagation algorithm to indicate how a machine should change its internal parameters that are used to compute the representation in each layer from the representation in the previous layer. Deep convolutional nets have brought about breakthroughs in processing images, video, speech and audio, whereas recurrent nets have shone light on sequential data such as text and speech.

Machine-learning technology powers many aspects of modern society: from web searches to content filtering on social networks to recommendations on e-commerce websites, and it is increasingly present in consumer products such as cameras and smartphones. Machine-learning systems are used to identify objects in images, transcribe speech into text, match news items, posts or products with users' interests, and select relevant results of search. Increasingly, these applications make use of a class of techniques called deep learning.

Conventional machine-learning techniques were limited in their ability to process natural data in their raw form. For decades, constructing a pattern-recognition or machine-learning system required careful engineering and considerable domain expertise to design a feature extractor that transformed the raw data (such as the pixel values of an image) into a suitable internal representation or feature vector from which the learning subsystem, often a classifier, could detect or classify patterns in the input.

Representation learning is a set of methods that allows a machine to be fed with raw data and to automatically discover the representations needed for detection or classification. Deep-learning methods are representation-learning methods with multiple levels of representation, obtained by composing simple but non-linear modules that each transform the representation at one level (starting with the raw input) into a representation at a higher, slightly more abstract level. With the composition of enough such transformations, very complex functions can be learned. For classification tasks, higher layers of representation amplify aspects of the input that are important for discrimination and suppress irrelevant variations. An image, for example, comes in the form of an array of pixel values, and the learned features in the first layer of representation typically represent the presence or absence of edges at particular orientations and locations in the image. The second layer typically detects motifs by spotting particular arrangements of edges, regardless of small variations in the edge positions. The third layer may assemble motifs into larger combinations that correspond to parts of familiar objects, and subsequent layers would detect objects as combinations of these parts. The key aspect of deep learning is that these layers of features are not designed by human engineers: they are learned from data using a general-purpose learning procedure.

Deep learning is making major advances in solving problems that have resisted the best attempts of the artificial intelligence community for many years. It has turned out to be very good at discovering

intricate structures in high-dimensional data and is therefore applicable to many domains of science, business and government. In addition to beating records in image recognition^{1–4} and speech recognition^{5–7}, it has beaten other machine-learning techniques at predicting the activity of potential drug molecules⁸, analysing particle accelerator data^{9,10}, reconstructing brain circuits¹¹, and predicting the effects of mutations in non-coding DNA on gene expression and disease^{12,13}. Perhaps more surprisingly, deep learning has produced extremely promising results for various tasks in natural language understanding¹⁴, particularly topic classification, sentiment analysis, question answering¹⁵ and language translation^{16,17}.

We think that deep learning will have many more successes in the near future because it requires very little engineering by hand, so it can easily take advantage of increases in the amount of available computation and data. New learning algorithms and architectures that are currently being developed for deep neural networks will only accelerate this progress.

Supervised learning

The most common form of machine learning, deep or not, is supervised learning. Imagine that we want to build a system that can classify images as containing, say, a house, a car, a person or a pet. We first collect a large data set of images of houses, cars, people and pets, each labelled with its category. During training, the machine is shown an image and produces an output in the form of a vector of scores, one for each category. We want the desired category to have the highest score of all categories, but this is unlikely to happen before training. We compute an objective function that measures the error (or distance) between the output scores and the desired pattern of scores. The machine then modifies its internal adjustable parameters to reduce this error. These adjustable parameters, often called weights, are real numbers that can be seen as 'knobs' that define the input–output function of the machine. In a typical deep-learning system, there may be hundreds of millions of these adjustable weights, and hundreds of millions of labelled examples with which to train the machine.

To properly adjust the weight vector, the learning algorithm computes a gradient vector that, for each weight, indicates by what amount the error would increase or decrease if the weight were increased by a tiny amount. The weight vector is then adjusted in the opposite direction to the gradient vector.

The objective function, averaged over all the training examples, can

¹Facebook AI Research, 770 Broadway, New York, New York 10003 USA. ²New York University, 715 Broadway, New York, New York 10003, USA. ³Department of Computer Science and Operations Research Université de Montréal, Pavillon André-Aisenstadt, PO Box 6128 Centre-Ville STN Montréal, Quebec H3C 3J7, Canada. ⁴Google, 1600 Amphitheatre Parkway, Mountain View, California 94043, USA. ⁵Department of Computer Science, University of Toronto, 6 King's College Road, Toronto, Ontario M5S 3G4, Canada.

be seen as a kind of hilly landscape in the high-dimensional space of weight values. The negative gradient vector indicates the direction of steepest descent in this landscape, taking it closer to a minimum, where the output error is low on average.

In practice, most practitioners use a procedure called stochastic gradient descent (SGD). This consists of showing the input vector for a few examples, computing the outputs and the errors, computing the average gradient for those examples, and adjusting the weights accordingly. The process is repeated for many small sets of examples from the training set until the average of the objective function stops decreasing. It is called stochastic because each small set of examples gives a noisy estimate of the average gradient over all examples. This simple procedure usually finds a good set of weights surprisingly quickly when compared with far more elaborate optimization techniques¹⁸. After training, the performance of the system is measured on a different set of examples called a test set. This serves to test the generalization ability of the machine — its ability to produce sensible answers on new inputs that it has never seen during training.

Many of the current practical applications of machine learning use linear classifiers on top of hand-engineered features. A two-class linear classifier computes a weighted sum of the feature vector components. If the weighted sum is above a threshold, the input is classified as belonging to a particular category.

Since the 1960s we have known that linear classifiers can only carve their input space into very simple regions, namely half-spaces separated by a hyperplane¹⁹. But problems such as image and speech recognition require the input–output function to be insensitive to irrelevant variations of the input, such as variations in position, orientation or illumination of an object, or variations in the pitch or accent of speech, while being very sensitive to particular minute variations (for example, the difference between a white wolf and a breed of wolf-like white dog called a Samoyed). At the pixel level, images of two Samoyeds in different poses and in different environments may be very different from each other, whereas two images of a Samoyed and a wolf in the same position and on similar backgrounds may be very similar to each other. A linear classifier, or any other ‘shallow’ classifier operating on

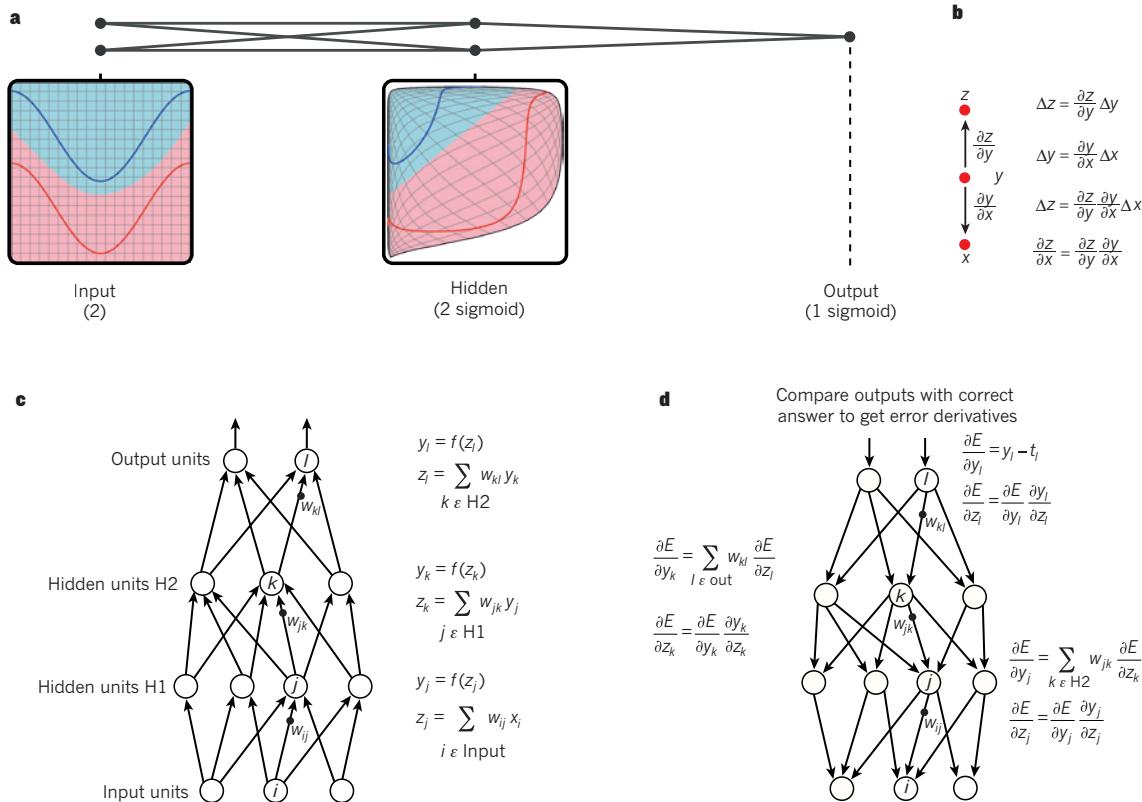


Figure 1 | Multilayer neural networks and backpropagation. **a**, A multilayer neural network (shown by the connected dots) can distort the input space to make the classes of data (examples of which are on the red and blue lines) linearly separable. Note how a regular grid (shown on the left) in input space is also transformed (shown in the middle panel) by hidden units. This is an illustrative example with only two input units, two hidden units and one output unit, but the networks used for object recognition or natural language processing contain tens or hundreds of thousands of units. Reproduced with permission from C. Olah (<http://colah.github.io/>). **b**, The chain rule of derivatives tells us how two small effects (that of a small change of x on y , and that of y on z) are composed. A small change Δx in x gets transformed first into a small change Δy in y by getting multiplied by $\partial y / \partial x$ (that is, the definition of partial derivative). Similarly, the change Δy creates a change Δz in z . Substituting one equation into the other gives the chain rule of derivatives — how Δx gets turned into Δz through multiplication by the product of $\partial y / \partial x$ and $\partial z / \partial y$. It also works when x , y and z are vectors (and the derivatives are Jacobian matrices). **c**, The equations used for computing the forward pass in a neural net with two hidden layers and one output layer, each constituting a module through

which one can backpropagate gradients. At each layer, we first compute the total input z to each unit, which is a weighted sum of the outputs of the units in the layer below. Then a non-linear function $f(\cdot)$ is applied to z to get the output of the unit. For simplicity, we have omitted bias terms. The non-linear functions used in neural networks include the rectified linear unit (ReLU) $f(z) = \max(0, z)$, commonly used in recent years, as well as the more conventional sigmoids, such as the hyperbolic tangent, $f(z) = (\exp(z) - \exp(-z)) / (\exp(z) + \exp(-z))$ and logistic function logistic, $f(z) = 1 / (1 + \exp(-z))$. **d**, The equations used for computing the backward pass. At each hidden layer we compute the error derivative with respect to the output of each unit, which is a weighted sum of the error derivatives with respect to the total inputs to the units in the layer above. We then convert the error derivative with respect to the input by multiplying it by the gradient of $f(z)$. At the output layer, the error derivative with respect to the output of a unit is computed by differentiating the cost function. This gives $y_l - t_l$ if the cost function for unit l is $0.5(y_l - t_l)^2$, where t_l is the target value. Once the $\partial E / \partial z$ is known, the error-derivative for the weight w_{jk} on the connection from unit j in the layer below is just $y_j \partial E / \partial z$.

Samoyed (16); Papillon (5.7); Pomeranian (2.7); Arctic fox (1.0); Eskimo dog (0.6); white wolf (0.4); Siberian husky (0.4)

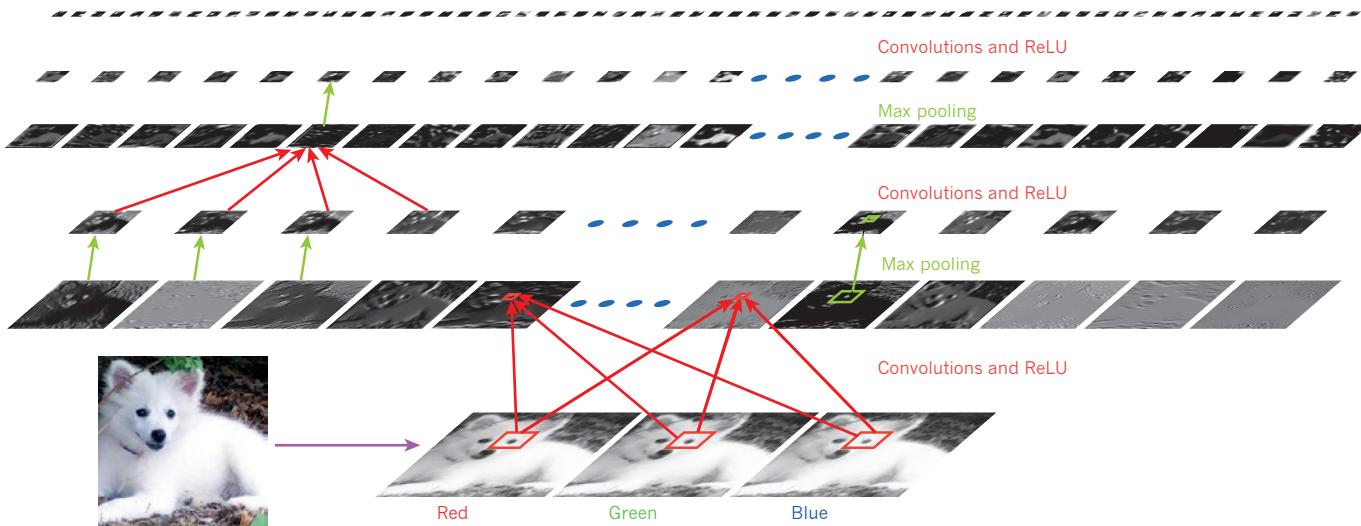


Figure 2 | Inside a convolutional network. The outputs (not the filters) of each layer (horizontally) of a typical convolutional network architecture applied to the image of a Samoyed dog (bottom left; and RGB (red, green, blue) inputs, bottom right). Each rectangular image is a feature map

raw pixels could not possibly distinguish the latter two, while putting the former two in the same category. This is why shallow classifiers require a good feature extractor that solves the selectivity–invariance dilemma — one that produces representations that are selective to the aspects of the image that are important for discrimination, but that are invariant to irrelevant aspects such as the pose of the animal. To make classifiers more powerful, one can use generic non-linear features, as with kernel methods²⁰, but generic features such as those arising with the Gaussian kernel do not allow the learner to generalize well far from the training examples²¹. The conventional option is to hand design good feature extractors, which requires a considerable amount of engineering skill and domain expertise. But this can all be avoided if good features can be learned automatically using a general-purpose learning procedure. This is the key advantage of deep learning.

A deep-learning architecture is a multilayer stack of simple modules, all (or most) of which are subject to learning, and many of which compute non-linear input–output mappings. Each module in the stack transforms its input to increase both the selectivity and the invariance of the representation. With multiple non-linear layers, say a depth of 5 to 20, a system can implement extremely intricate functions of its inputs that are simultaneously sensitive to minute details — distinguishing Samoyeds from white wolves — and insensitive to large irrelevant variations such as the background, pose, lighting and surrounding objects.

Backpropagation to train multilayer architectures

From the earliest days of pattern recognition^{22,23}, the aim of researchers has been to replace hand-engineered features with trainable multilayer networks, but despite its simplicity, the solution was not widely understood until the mid 1980s. As it turns out, multilayer architectures can be trained by simple stochastic gradient descent. As long as the modules are relatively smooth functions of their inputs and of their internal weights, one can compute gradients using the backpropagation procedure. The idea that this could be done, and that it worked, was discovered independently by several different groups during the 1970s and 1980s^{24–27}.

The backpropagation procedure to compute the gradient of an objective function with respect to the weights of a multilayer stack of modules is nothing more than a practical application of the chain

corresponding to the output for one of the learned features, detected at each of the image positions. Information flows bottom up, with lower-level features acting as oriented edge detectors, and a score is computed for each image class in output. ReLU, rectified linear unit.

rule for derivatives. The key insight is that the derivative (or gradient) of the objective with respect to the input of a module can be computed by working backwards from the gradient with respect to the output of that module (or the input of the subsequent module) (Fig. 1). The backpropagation equation can be applied repeatedly to propagate gradients through all modules, starting from the output at the top (where the network produces its prediction) all the way to the bottom (where the external input is fed). Once these gradients have been computed, it is straightforward to compute the gradients with respect to the weights of each module.

Many applications of deep learning use feedforward neural network architectures (Fig. 1), which learn to map a fixed-size input (for example, an image) to a fixed-size output (for example, a probability for each of several categories). To go from one layer to the next, a set of units compute a weighted sum of their inputs from the previous layer and pass the result through a non-linear function. At present, the most popular non-linear function is the rectified linear unit (ReLU), which is simply the half-wave rectifier $f(z) = \max(z, 0)$. In past decades, neural nets used smoother non-linearities, such as $\tanh(z)$ or $1/(1 + \exp(-z))$, but the ReLU typically learns much faster in networks with many layers, allowing training of a deep supervised network without unsupervised pre-training²⁸. Units that are not in the input or output layer are conventionally called hidden units. The hidden layers can be seen as distorting the input in a non-linear way so that categories become linearly separable by the last layer (Fig. 1).

In the late 1990s, neural nets and backpropagation were largely forsaken by the machine-learning community and ignored by the computer-vision and speech-recognition communities. It was widely thought that learning useful, multistage, feature extractors with little prior knowledge was infeasible. In particular, it was commonly thought that simple gradient descent would get trapped in poor local minima — weight configurations for which no small change would reduce the average error.

In practice, poor local minima are rarely a problem with large networks. Regardless of the initial conditions, the system nearly always reaches solutions of very similar quality. Recent theoretical and empirical results strongly suggest that local minima are not a serious issue in general. Instead, the landscape is packed with a combinatorially large number of saddle points where the gradient is zero, and the surface curves up in most dimensions and curves down in the

remainder^{29,30}. The analysis seems to show that saddle points with only a few downward curving directions are present in very large numbers, but almost all of them have very similar values of the objective function. Hence, it does not much matter which of these saddle points the algorithm gets stuck at.

Interest in deep feedforward networks was revived around 2006 (refs 31–34) by a group of researchers brought together by the Canadian Institute for Advanced Research (CIFAR). The researchers introduced unsupervised learning procedures that could create layers of feature detectors without requiring labelled data. The objective in learning each layer of feature detectors was to be able to reconstruct or model the activities of feature detectors (or raw inputs) in the layer below. By ‘pre-training’ several layers of progressively more complex feature detectors using this reconstruction objective, the weights of a deep network could be initialized to sensible values. A final layer of output units could then be added to the top of the network and the whole deep system could be fine-tuned using standard backpropagation^{33–35}. This worked remarkably well for recognizing handwritten digits or for detecting pedestrians, especially when the amount of labelled data was very limited³⁶.

The first major application of this pre-training approach was in speech recognition, and it was made possible by the advent of fast graphics processing units (GPUs) that were convenient to program³⁷ and allowed researchers to train networks 10 or 20 times faster. In 2009, the approach was used to map short temporal windows of coefficients extracted from a sound wave to a set of probabilities for the various fragments of speech that might be represented by the frame in the centre of the window. It achieved record-breaking results on a standard speech recognition benchmark that used a small vocabulary³⁸ and was quickly developed to give record-breaking results on a large vocabulary task³⁹. By 2012, versions of the deep net from 2009 were being developed by many of the major speech groups⁶ and were already being deployed in Android phones. For smaller data sets, unsupervised pre-training helps to prevent overfitting⁴⁰, leading to significantly better generalization when the number of labelled examples is small, or in a transfer setting where we have lots of examples for some ‘source’ tasks but very few for some ‘target’ tasks. Once deep learning had been rehabilitated, it turned out that the pre-training stage was only needed for small data sets.

There was, however, one particular type of deep, feedforward network that was much easier to train and generalized much better than networks with full connectivity between adjacent layers. This was the convolutional neural network (ConvNet)^{41,42}. It achieved many practical successes during the period when neural networks were out of favour and it has recently been widely adopted by the computer-vision community.

Convolutional neural networks

ConvNets are designed to process data that come in the form of multiple arrays, for example a colour image composed of three 2D arrays containing pixel intensities in the three colour channels. Many data modalities are in the form of multiple arrays: 1D for signals and sequences, including language; 2D for images or audio spectrograms; and 3D for video or volumetric images. There are four key ideas behind ConvNets that take advantage of the properties of natural signals: local connections, shared weights, pooling and the use of many layers.

The architecture of a typical ConvNet (Fig. 2) is structured as a series of stages. The first few stages are composed of two types of layers: convolutional layers and pooling layers. Units in a convolutional layer are organized in feature maps, within which each unit is connected to local patches in the feature maps of the previous layer through a set of weights called a filter bank. The result of this local weighted sum is then passed through a non-linearity such as a ReLU. All units in a feature map share the same filter bank. Different feature maps in a layer use different filter banks. The reason for

this architecture is twofold. First, in array data such as images, local groups of values are often highly correlated, forming distinctive local motifs that are easily detected. Second, the local statistics of images and other signals are invariant to location. In other words, if a motif can appear in one part of the image, it could appear anywhere, hence the idea of units at different locations sharing the same weights and detecting the same pattern in different parts of the array. Mathematically, the filtering operation performed by a feature map is a discrete convolution, hence the name.

Although the role of the convolutional layer is to detect local conjunctions of features from the previous layer, the role of the pooling layer is to merge semantically similar features into one. Because the relative positions of the features forming a motif can vary somewhat, reliably detecting the motif can be done by coarse-graining the position of each feature. A typical pooling unit computes the maximum of a local patch of units in one feature map (or in a few feature maps). Neighbouring pooling units take input from patches that are shifted by more than one row or column, thereby reducing the dimension of the representation and creating an invariance to small shifts and distortions. Two or three stages of convolution, non-linearity and pooling are stacked, followed by more convolutional and fully-connected layers. Backpropagating gradients through a ConvNet is as simple as through a regular deep network, allowing all the weights in all the filter banks to be trained.

Deep neural networks exploit the property that many natural signals are compositional hierarchies, in which higher-level features are obtained by composing lower-level ones. In images, local combinations of edges form motifs, motifs assemble into parts, and parts form objects. Similar hierarchies exist in speech and text from sounds to phones, phonemes, syllables, words and sentences. The pooling allows representations to vary very little when elements in the previous layer vary in position and appearance.

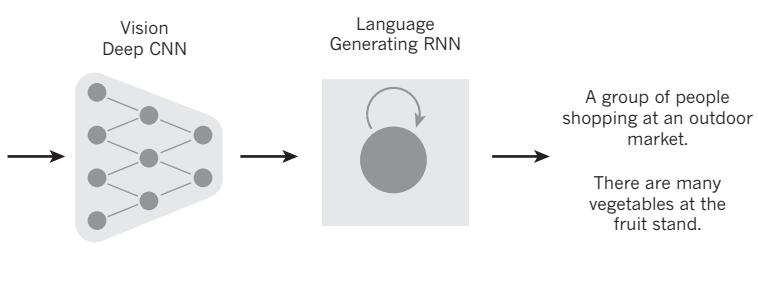
The convolutional and pooling layers in ConvNets are directly inspired by the classic notions of simple cells and complex cells in visual neuroscience⁴³, and the overall architecture is reminiscent of the LGN–V1–V2–V4–IT hierarchy in the visual cortex ventral pathway⁴⁴. When ConvNet models and monkeys are shown the same picture, the activations of high-level units in the ConvNet explains half of the variance of random sets of 160 neurons in the monkey’s inferotemporal cortex⁴⁵. ConvNets have their roots in the neocognitron⁴⁶, the architecture of which was somewhat similar, but did not have an end-to-end supervised-learning algorithm such as backpropagation. A primitive 1D ConvNet called a time-delay neural net was used for the recognition of phonemes and simple words^{47,48}.

There have been numerous applications of convolutional networks going back to the early 1990s, starting with time-delay neural networks for speech recognition⁴⁷ and document reading⁴². The document reading system used a ConvNet trained jointly with a probabilistic model that implemented language constraints. By the late 1990s this system was reading over 10% of all the cheques in the United States. A number of ConvNet-based optical character recognition and handwriting recognition systems were later deployed by Microsoft⁴⁹. ConvNets were also experimented with in the early 1990s for object detection in natural images, including faces and hands^{50,51}, and for face recognition⁵².

Image understanding with deep convolutional networks

Since the early 2000s, ConvNets have been applied with great success to the detection, segmentation and recognition of objects and regions in images. These were all tasks in which labelled data was relatively abundant, such as traffic sign recognition⁵³, the segmentation of biological images⁵⁴ particularly for connectomics⁵⁵, and the detection of faces, text, pedestrians and human bodies in natural images^{36,50,51,56–58}. A major recent practical success of ConvNets is face recognition⁵⁹.

Importantly, images can be labelled at the pixel level, which will have applications in technology, including autonomous mobile robots and



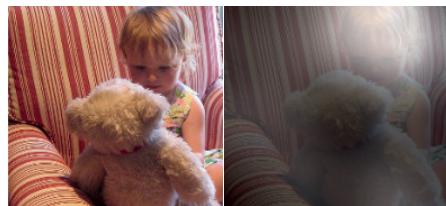
A woman is throwing a **frisbee** in a park.



A **dog** is standing on a hardwood floor.



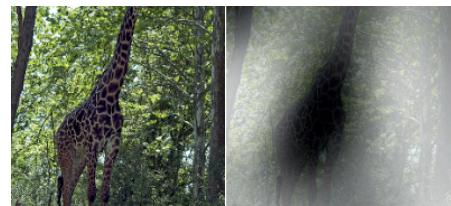
A **stop** sign is on a road with a mountain in the background



A little **girl** sitting on a bed with a **teddy bear**.



A group of **people** sitting on a boat in the water.



A **giraffe** standing in a forest with **trees** in the background.

Figure 3 | From image to text. Captions generated by a recurrent neural network (RNN) taking, as extra input, the representation extracted by a deep convolutional neural network (CNN) from a test image, with the RNN trained to ‘translate’ high-level representations of images into captions (top). Reproduced

self-driving cars^{60,61}. Companies such as Mobileye and NVIDIA are using such ConvNet-based methods in their upcoming vision systems for cars. Other applications gaining importance involve natural language understanding¹⁴ and speech recognition⁷.

Despite these successes, ConvNets were largely forsaken by the mainstream computer-vision and machine-learning communities until the ImageNet competition in 2012. When deep convolutional networks were applied to a data set of about a million images from the web that contained 1,000 different classes, they achieved spectacular results, almost halving the error rates of the best competing approaches¹. This success came from the efficient use of GPUs, ReLUs, a new regularization technique called dropout⁶², and techniques to generate more training examples by deforming the existing ones. This success has brought about a revolution in computer vision; ConvNets are now the dominant approach for almost all recognition and detection tasks^{4,58,59,63–65} and approach human performance on some tasks. A recent stunning demonstration combines ConvNets and recurrent net modules for the generation of image captions (Fig. 3).

Recent ConvNet architectures have 10 to 20 layers of ReLUs, hundreds of millions of weights, and billions of connections between units. Whereas training such large networks could have taken weeks only two years ago, progress in hardware, software and algorithm parallelization have reduced training times to a few hours.

The performance of ConvNet-based vision systems has caused most major technology companies, including Google, Facebook,

with permission from ref. 102. When the RNN is given the ability to focus its attention on a different location in the input image (middle and bottom; the lighter patches were given more attention) as it generates each word (bold), we found⁶⁶ that it exploits this to achieve better ‘translation’ of images into captions.

Microsoft, IBM, Yahoo!, Twitter and Adobe, as well as a quickly growing number of start-ups to initiate research and development projects and to deploy ConvNet-based image understanding products and services.

ConvNets are easily amenable to efficient hardware implementations in chips or field-programmable gate arrays^{66,67}. A number of companies such as NVIDIA, Mobileye, Intel, Qualcomm and Samsung are developing ConvNet chips to enable real-time vision applications in smartphones, cameras, robots and self-driving cars.

Distributed representations and language processing

Deep-learning theory shows that deep nets have two different exponential advantages over classic learning algorithms that do not use distributed representations²¹. Both of these advantages arise from the power of composition and depend on the underlying data-generating distribution having an appropriate componential structure⁴⁰. First, learning distributed representations enable generalization to new combinations of the values of learned features beyond those seen during training (for example, 2^n combinations are possible with n binary features)^{68,69}. Second, composing layers of representation in a deep net brings the potential for another exponential advantage⁷⁰ (exponential in the depth).

The hidden layers of a multilayer neural network learn to represent the network’s inputs in a way that makes it easy to predict the target outputs. This is nicely demonstrated by training a multilayer neural network to predict the next word in a sequence from a local

context of earlier words⁷¹. Each word in the context is presented to the network as a one-of-N vector, that is, one component has a value of 1 and the rest are 0. In the first layer, each word creates a different pattern of activations, or word vectors (Fig. 4). In a language model, the other layers of the network learn to convert the input word vectors into an output word vector for the predicted next word, which can be used to predict the probability for any word in the vocabulary to appear as the next word. The network learns word vectors that contain many active components each of which can be interpreted as a separate feature of the word, as was first demonstrated²⁷ in the context of learning distributed representations for symbols. These semantic features were not explicitly present in the input. They were discovered by the learning procedure as a good way of factorizing the structured relationships between the input and output symbols into multiple ‘micro-rules’. Learning word vectors turned out to also work very well when the word sequences come from a large corpus of real text and the individual micro-rules are unreliable⁷¹. When trained to predict the next word in a news story, for example, the learned word vectors for Tuesday and Wednesday are very similar, as are the word vectors for Sweden and Norway. Such representations are called distributed representations because their elements (the features) are not mutually exclusive and their many configurations correspond to the variations seen in the observed data. These word vectors are composed of learned features that were not determined ahead of time by experts, but automatically discovered by the neural network. Vector representations of words learned from text are now very widely used in natural language applications^{14,17,72–76}.

The issue of representation lies at the heart of the debate between the logic-inspired and the neural-network-inspired paradigms for cognition. In the logic-inspired paradigm, an instance of a symbol is something for which the only property is that it is either identical or non-identical to other symbol instances. It has no internal structure that is relevant to its use; and to reason with symbols, they must be bound to the variables in judiciously chosen rules of inference. By contrast, neural networks just use big activity vectors, big weight matrices and scalar non-linearities to perform the type of fast ‘intuitive’ inference that underpins effortless commonsense reasoning.

Before the introduction of neural language models⁷¹, the standard approach to statistical modelling of language did not exploit distributed representations: it was based on counting frequencies of occurrences of short symbol sequences of length up to N (called N-grams). The number of possible N-grams is on the order of V^N , where V is the vocabulary size, so taking into account a context of more than a

handful of words would require very large training corpora. N-grams treat each word as an atomic unit, so they cannot generalize across semantically related sequences of words, whereas neural language models can because they associate each word with a vector of real valued features, and semantically related words end up close to each other in that vector space (Fig. 4).

Recurrent neural networks

When backpropagation was first introduced, its most exciting use was for training recurrent neural networks (RNNs). For tasks that involve sequential inputs, such as speech and language, it is often better to use RNNs (Fig. 5). RNNs process an input sequence one element at a time, maintaining in their hidden units a ‘state vector’ that implicitly contains information about the history of all the past elements of the sequence. When we consider the outputs of the hidden units at different discrete time steps as if they were the outputs of different neurons in a deep multilayer network (Fig. 5, right), it becomes clear how we can apply backpropagation to train RNNs.

RNNs are very powerful dynamic systems, but training them has proved to be problematic because the backpropagated gradients either grow or shrink at each time step, so over many time steps they typically explode or vanish^{77,78}.

Thanks to advances in their architecture^{79,80} and ways of training them^{81,82}, RNNs have been found to be very good at predicting the next character in the text⁸³ or the next word in a sequence⁷⁵, but they can also be used for more complex tasks. For example, after reading an English sentence one word at a time, an English ‘encoder’ network can be trained so that the final state vector of its hidden units is a good representation of the thought expressed by the sentence. This thought vector can then be used as the initial hidden state of (or as extra input to) a jointly trained French ‘decoder’ network, which outputs a probability distribution for the first word of the French translation. If a particular first word is chosen from this distribution and provided as input to the decoder network it will then output a probability distribution for the second word of the translation and so on until a full stop is chosen^{17,72,76}. Overall, this process generates sequences of French words according to a probability distribution that depends on the English sentence. This rather naive way of performing machine translation has quickly become competitive with the state-of-the-art, and this raises serious doubts about whether understanding a sentence requires anything like the internal symbolic expressions that are manipulated by using inference rules. It is more compatible with the view that everyday reasoning involves many simultaneous analogies

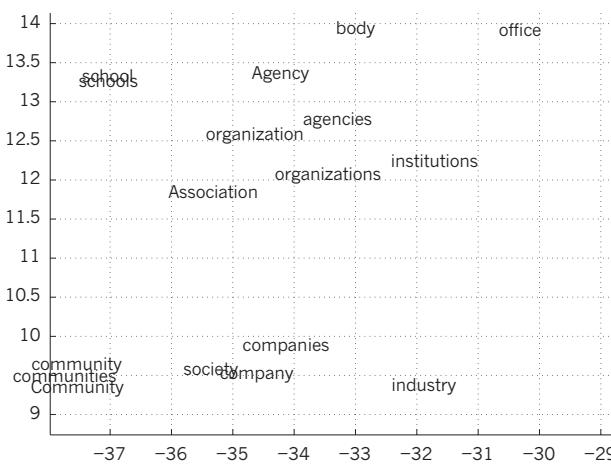


Figure 4 | Visualizing the learned word vectors. On the left is an illustration of word representations learned for modelling language, non-linearly projected to 2D for visualization using the t-SNE algorithm¹⁰³. On the right is a 2D representation of phrases learned by an English-to-French encoder-decoder recurrent neural network⁷⁵. One can observe that semantically similar words

or sequences of words are mapped to nearby representations. The distributed representations of words are obtained by using backpropagation to jointly learn a representation for each word and a function that predicts a target quantity such as the next word in a sequence (for language modelling) or a whole sequence of translated words (for machine translation)^{18,75}.

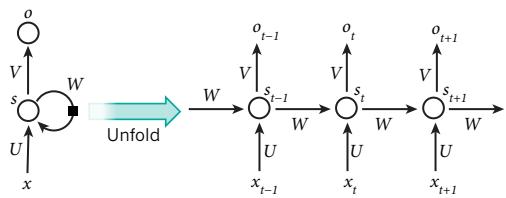


Figure 5 | A recurrent neural network and the unfolding in time of the computation involved in its forward computation. The artificial neurons (for example, hidden units grouped under node s with values s_t at time t) get inputs from other neurons at previous time steps (this is represented with the black square, representing a delay of one time step, on the left). In this way, a recurrent neural network can map an input sequence with elements x_t into an output sequence with elements o_t , with each o_t depending on all the previous x_t' (for $t' \leq t$). The same parameters (matrices U, V, W) are used at each time step. Many other architectures are possible, including a variant in which the network can generate a sequence of outputs (for example, words), each of which is used as inputs for the next time step. The backpropagation algorithm (Fig. 1) can be directly applied to the computational graph of the unfolded network on the right, to compute the derivative of a total error (for example, the log-probability of generating the right sequence of outputs) with respect to all the states s , and all the parameters.

that each contribute plausibility to a conclusion^{84,85}.

Instead of translating the meaning of a French sentence into an English sentence, one can learn to ‘translate’ the meaning of an image into an English sentence (Fig. 3). The encoder here is a deep ConvNet that converts the pixels into an activity vector in its last hidden layer. The decoder is an RNN similar to the ones used for machine translation and neural language modelling. There has been a surge of interest in such systems recently (see examples mentioned in ref. 86).

RNNs, once unfolded in time (Fig. 5), can be seen as very deep feedforward networks in which all the layers share the same weights. Although their main purpose is to learn long-term dependencies, theoretical and empirical evidence shows that it is difficult to learn to store information for very long⁷⁸.

To correct for that, one idea is to augment the network with an explicit memory. The first proposal of this kind is the long short-term memory (LSTM) networks that use special hidden units, the natural behaviour of which is to remember inputs for a long time⁷⁹. A special unit called the memory cell acts like an accumulator or a gated leaky neuron: it has a connection to itself at the next time step that has a weight of one, so it copies its own real-valued state and accumulates the external signal, but this self-connection is multiplicatively gated by another unit that learns to decide when to clear the content of the memory.

LSTM networks have subsequently proved to be more effective than conventional RNNs, especially when they have several layers for each time step⁸⁷, enabling an entire speech recognition system that goes all the way from acoustics to the sequence of characters in the transcription. LSTM networks or related forms of gated units are also currently used for the encoder and decoder networks that perform so well at machine translation^{17,72,76}.

Over the past year, several authors have made different proposals to augment RNNs with a memory module. Proposals include the Neural Turing Machine in which the network is augmented by a ‘tape-like’ memory that the RNN can choose to read from or write to⁸⁸, and memory networks, in which a regular network is augmented by a kind of associative memory⁸⁹. Memory networks have yielded excellent performance on standard question-answering benchmarks. The memory is used to remember the story about which the network is later asked to answer questions.

Beyond simple memorization, neural Turing machines and memory networks are being used for tasks that would normally require reasoning and symbol manipulation. Neural Turing machines can be taught ‘algorithms’. Among other things, they can learn to output

a sorted list of symbols when their input consists of an unsorted sequence in which each symbol is accompanied by a real value that indicates its priority in the list⁸⁸. Memory networks can be trained to keep track of the state of the world in a setting similar to a text adventure game and after reading a story, they can answer questions that require complex inference⁹⁰. In one test example, the network is shown a 15-sentence version of the *The Lord of the Rings* and correctly answers questions such as “where is Frodo now?”⁸⁹.

The future of deep learning

Unsupervised learning^{91–98} had a catalytic effect in reviving interest in deep learning, but has since been overshadowed by the successes of purely supervised learning. Although we have not focused on it in this Review, we expect unsupervised learning to become far more important in the longer term. Human and animal learning is largely unsupervised: we discover the structure of the world by observing it, not by being told the name of every object.

Human vision is an active process that sequentially samples the optic array in an intelligent, task-specific way using a small, high-resolution fovea with a large, low-resolution surround. We expect much of the future progress in vision to come from systems that are trained end-to-end and combine ConvNets with RNNs that use reinforcement learning to decide where to look. Systems combining deep learning and reinforcement learning are in their infancy, but they already outperform passive vision systems⁹⁹ at classification tasks and produce impressive results in learning to play many different video games¹⁰⁰.

Natural language understanding is another area in which deep learning is poised to make a large impact over the next few years. We expect systems that use RNNs to understand sentences or whole documents will become much better when they learn strategies for selectively attending to one part at a time^{76,86}.

Ultimately, major progress in artificial intelligence will come about through systems that combine representation learning with complex reasoning. Although deep learning and simple reasoning have been used for speech and handwriting recognition for a long time, new paradigms are needed to replace rule-based manipulation of symbolic expressions by operations on large vectors¹⁰¹. ■

Received 25 February; accepted 1 May 2015.

- Krizhevsky, A., Sutskever, I. & Hinton, G. ImageNet classification with deep convolutional neural networks. In *Proc. Advances in Neural Information Processing Systems 25* 1090–1098 (2012). **This report was a breakthrough that used convolutional nets to almost halve the error rate for object recognition, and precipitated the rapid adoption of deep learning by the computer vision community.**
- Farabet, C., Couprie, C., Najman, L. & LeCun, Y. Learning hierarchical features for scene labeling. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**, 1915–1929 (2013).
- Tompson, J., Jain, A., LeCun, Y. & Bregler, C. Joint training of a convolutional network and a graphical model for human pose estimation. In *Proc. Advances in Neural Information Processing Systems 27* 1799–1807 (2014).
- Szegedy, C. et al. Going deeper with convolutions. Preprint at <http://arxiv.org/abs/1409.4842> (2014).
- Mikolov, T., Deoras, A., Povey, D., Burget, L. & Cernocky, J. Strategies for training large scale neural network language models. In *Proc. Automatic Speech Recognition and Understanding* 196–201 (2011).
- Hinton, G. et al. Deep neural networks for acoustic modeling in speech recognition. *IEEE Signal Processing Magazine* **29**, 82–97 (2012). **This joint paper from the major speech recognition laboratories, summarizing the breakthrough achieved with deep learning on the task of phonetic classification for automatic speech recognition, was the first major industrial application of deep learning.**
- Sainath, T., Mohamed, A.-R., Kingsbury, B. & Ramabhadran, B. Deep convolutional neural networks for LVCSR. In *Proc. Acoustics, Speech and Signal Processing* 8614–8618 (2013).
- Ma, J., Sheridan, R. P., Liaw, A., Dahl, G. E. & Svetnik, V. Deep neural nets as a method for quantitative structure-activity relationships. *J. Chem. Inf. Model.* **55**, 263–274 (2015).
- Ciodaro, T., Deva, D., de Seixas, J. & Damazio, D. Online particle detection with neural networks based on topological calorimetry information. *J. Phys. Conf. Series* **368**, 012030 (2012).
- Kaggle. Higgs boson machine learning challenge. Kaggle <https://www.kaggle.com/c/higgs-boson> (2014).
- Helmstaedter, M. et al. Connectomic reconstruction of the inner plexiform layer in the mouse retina. *Nature* **500**, 168–174 (2013).

12. Leung, M. K., Xiong, H. Y., Lee, L. J. & Frey, B. J. Deep learning of the tissue-regulated splicing code. *Bioinformatics* **30**, i121–i129 (2014).

13. Xiong, H. Y. et al. The human splicing code reveals new insights into the genetic determinants of disease. *Science* **347**, 6218 (2015).

14. Collobert, R., et al. Natural language processing (almost) from scratch. *J. Mach. Learn. Res.* **12**, 2493–2537 (2011).

15. Bordes, A., Chopra, S. & Weston, J. Question answering with subgraph embeddings. In *Proc. Empirical Methods in Natural Language Processing* <http://arxiv.org/abs/1406.3676v3> (2014).

16. Jean, S., Cho, K., Memisevic, R. & Bengio, Y. On using very large target vocabulary for neural machine translation. In *Proc. ACL-IJCNLP* <http://arxiv.org/abs/1412.2007> (2015).

17. Sutskever, I., Vinyals, O. & Le, Q. V. Sequence to sequence learning with neural networks. In *Proc. Advances in Neural Information Processing Systems* **27** 3104–3112 (2014). **This paper showed state-of-the-art machine translation results with the architecture introduced in ref. 72, with a recurrent network trained to read a sentence in one language, produce a semantic representation of its meaning, and generate a translation in another language.**

18. Bottou, L. & Bousquet, O. The tradeoffs of large scale learning. In *Proc. Advances in Neural Information Processing Systems* **20** 161–168 (2007).

19. Duda, R. O. & Hart, P. E. *Pattern Classification and Scene Analysis* (Wiley, 1973).

20. Schölkopf, B. & Smola, A. *Learning with Kernels* (MIT Press, 2002).

21. Bengio, Y., Delalleau, O. & Le Roux, N. The curse of highly variable functions for local kernel machines. In *Proc. Advances in Neural Information Processing Systems* **18** 107–114 (2005).

22. Selfridge, O. G. Pandemonium: a paradigm for learning in mechanisation of thought processes. In *Proc. Symposium on Mechanisation of Thought Processes* 513–526 (1958).

23. Rosenblatt, F. *The Perceptron — A Perceiving and Recognizing Automaton*. Tech. Rep. 85-460-1 (Cornell Aeronautical Laboratory, 1957).

24. Werbos, P. *Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences*. PhD thesis, Harvard Univ. (1974).

25. Parker, D. B. *Learning Logic Report TR-47* (MIT Press, 1985).

26. LeCun, Y. Une procédure d'apprentissage pour Réseau à seuil assymétrique in *Cognitiva '85: à la Frontière de l'Intelligence Artificielle, des Sciences de la Connaissance et des Neurosciences* [in French] 599–604 (1985).

27. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. *Nature* **323**, 533–536 (1986).

28. Glorot, X., Bordes, A. & Bengio, Y. Deep sparse rectifier neural networks. In *Proc. 14th International Conference on Artificial Intelligence and Statistics* 315–323 (2011). **This paper showed that supervised training of very deep neural networks is much faster if the hidden layers are composed of ReLU.**

29. Dauphin, Y. et al. Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. In *Proc. Advances in Neural Information Processing Systems* **27** 2933–2941 (2014).

30. Choromanska, A., Henaff, M., Mathieu, M., Arous, G. B. & LeCun, Y. The loss surface of multilayer networks. In *Proc. Conference on AI and Statistics* <http://arxiv.org/abs/1412.0233> (2014).

31. Hinton, G. E. What kind of graphical model is the brain? In *Proc. 19th International Joint Conference on Artificial intelligence* 1765–1775 (2005).

32. Hinton, G. E., Osindero, S. & Teh, Y.-W. A fast learning algorithm for deep belief nets. *Neural Comp.* **18**, 1527–1554 (2006). **This paper introduced a novel and effective way of training very deep neural networks by pre-training one hidden layer at a time using the unsupervised learning procedure for restricted Boltzmann machines.**

33. Bengio, Y., Lamblin, P., Popovici, D. & Larochelle, H. Greedy layer-wise training of deep networks. In *Proc. Advances in Neural Information Processing Systems* **19** 153–160 (2006). **This report demonstrated that the unsupervised pre-training method introduced in ref. 32 significantly improves performance on test data and generalizes the method to other unsupervised representation-learning techniques, such as auto-encoders.**

34. Ranzato, M., Poultney, C., Chopra, S. & LeCun, Y. Efficient learning of sparse representations with an energy-based model. In *Proc. Advances in Neural Information Processing Systems* **19** 1137–1144 (2006).

35. Hinton, G. E. & Salakhutdinov, R. Reducing the dimensionality of data with neural networks. *Science* **313**, 504–507 (2006).

36. Sermanet, P., Kavukcuoglu, K., Chintala, S. & LeCun, Y. Pedestrian detection with unsupervised multi-stage feature learning. In *Proc. International Conference on Computer Vision and Pattern Recognition* <http://arxiv.org/abs/1212.0142> (2013).

37. Raina, R., Madhavan, A. & Ng, A. Y. Large-scale deep unsupervised learning using graphics processors. In *Proc. 26th Annual International Conference on Machine Learning* 873–880 (2009).

38. Mohamed, A.-R., Dahl, G. E. & Hinton, G. Acoustic modeling using deep belief networks. *IEEE Trans. Audio Speech Lang. Process.* **20**, 14–22 (2012).

39. Dahl, G. E., Yu, D., Deng, L. & Acero, A. Context-dependent pre-trained deep neural networks for large vocabulary speech recognition. *IEEE Trans. Audio Speech Lang. Process.* **20**, 33–42 (2012).

40. Bengio, Y., Courville, A. & Vincent, P. Representation learning: a review and new perspectives. *IEEE Trans. Pattern Anal. Machine Intell.* **35**, 1798–1828 (2013). **This paper introduced neural language models, which learn to convert a word symbol into a word vector or word embedding composed of learned semantic features in order to predict the next word in a sequence.**

41. LeCun, Y. et al. Handwritten digit recognition with a back-propagation network. In *Proc. Advances in Neural Information Processing Systems* 396–404 (1990). **This is the first paper on convolutional networks trained by backpropagation for the task of classifying low-resolution images of handwritten digits.**

42. LeCun, Y., Bottou, L., Bengio, Y. & Haffner, P. Gradient-based learning applied to document recognition. *Proc. IEEE* **86**, 2278–2324 (1998). **This overview paper on the principles of end-to-end training of modular systems such as deep neural networks using gradient-based optimization showed how neural networks (and in particular convolutional nets) can be combined with search or inference mechanisms to model complex outputs that are interdependent, such as sequences of characters associated with the content of a document.**

43. Hubel, D. H. & Wiesel, T. N. Receptive fields, binocular interaction, and functional architecture in the cat's visual cortex. *J. Physiol.* **160**, 106–154 (1962).

44. Fellman, D. J. & Essen, D. C. V. Distributed hierarchical processing in the primate cerebral cortex. *Cereb. Cortex* **1**, 1–47 (1991).

45. Cadieu, C. F. et al. Deep neural networks rival the representation of primate it cortex for core visual object recognition. *PLoS Comp. Biol.* **10**, e1003963 (2014).

46. Fukushima, K. & Miyake, S. Neocognitron: a new algorithm for pattern recognition tolerant of deformations and shifts in position. *Pattern Recognition* **15**, 455–469 (1982).

47. Waibel, A., Hanazawa, T., Hinton, G. E., Shikano, K. & Lang, K. Phoneme recognition using time-delay neural networks. *IEEE Trans. Acoustics Speech Signal Process.* **37**, 328–339 (1989).

48. Bottou, L., Fogelman-Soulie, F., Blanchet, P. & Lienard, J. Experiments with time delay networks and dynamic time warping for speaker independent isolated digit recognition. In *Proc. EuroSpeech '89* 537–540 (1989).

49. Simard, D., Steinhaus, P. Y. & Platt, J. C. Best practices for convolutional neural networks. In *Proc. Document Analysis and Recognition* 958–963 (2003).

50. Vaillant, R., Monrocq, C. & LeCun, Y. Original approach for the localisation of objects in images. In *Proc. Vision, Image, and Signal Processing* **141**, 245–250 (1994).

51. Nowlan, S. & Platt, J. In *Neural Information Processing Systems* 901–908 (1995).

52. Lawrence, S., Giles, C. L., Tsoi, A. C. & Back, A. D. Face recognition: a convolutional neural-network approach. *IEEE Trans. Neural Networks* **8**, 98–113 (1997).

53. Ciresan, D., Meier, U., Masci, J. & Schmidhuber, J. Multi-column deep neural network for traffic sign classification. *Neural Networks* **32**, 333–338 (2012).

54. Ning, F. et al. Toward automatic phenotyping of developing embryos from videos. *IEEE Trans. Image Process.* **14**, 1360–1371 (2005).

55. Turaga, S. C. et al. Convolutional networks can learn to generate affinity graphs for image segmentation. *Neural Comput.* **22**, 511–538 (2010).

56. Garcia, C. & Delakis, M. Convolutional face finder: a neural architecture for fast and robust face detection. *IEEE Trans. Pattern Anal. Machine Intell.* **26**, 1408–1423 (2004).

57. Osadchy, M., LeCun, Y. & Miller, M. Synergistic face detection and pose estimation with energy-based models. *J. Mach. Learn. Res.* **8**, 1197–1215 (2007).

58. Tompson, J., Goroshin, R. R., Jain, A., LeCun, Y. Y. & Bregler, C. C. Efficient object localization using convolutional networks. In *Proc. Conference on Computer Vision and Pattern Recognition* <http://arxiv.org/abs/1411.4280> (2014).

59. Taigman, Y., Yang, M., Ranzato, M. & Wolf, L. Deepface: closing the gap to human-level performance in face verification. In *Proc. Conference on Computer Vision and Pattern Recognition* 1701–1708 (2014).

60. Hadsell, R. et al. Learning long-range vision for autonomous off-road driving. *J. Field Robot.* **26**, 120–144 (2009).

61. Farabet, C., Courville, C., Najman, L. & LeCun, Y. Scene parsing with multiscale feature learning, purity trees, and optimal covers. In *Proc. International Conference on Machine Learning* <http://arxiv.org/abs/1202.2160> (2012).

62. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I. & Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *J. Machine Learning Res.* **15**, 1929–1958 (2014).

63. Sermanet, P. et al. Overfeat: integrated recognition, localization and detection using convolutional networks. In *Proc. International Conference on Learning Representations* <http://arxiv.org/abs/1312.6229> (2014).

64. Girshick, R., Donahue, J., Darrell, T. & Malik, J. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proc. Conference on Computer Vision and Pattern Recognition* 580–587 (2014).

65. Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. In *Proc. International Conference on Learning Representations* <http://arxiv.org/abs/1409.1556> (2014).

66. Boser, B., Sackinger, E., Bromley, J., LeCun, Y. & Jackel, L. An analog neural network processor with programmable topology. *J. Solid State Circuits* **26**, 2017–2025 (1991).

67. Farabet, C. et al. Large-scale FPGA-based convolutional networks. In *Scaling up Machine Learning: Parallel and Distributed Approaches* (eds Bekkerman, R., Bilenko, M. & Langford, J.) 399–419 (Cambridge Univ. Press, 2011).

68. Bengio, Y. *Learning Deep Architectures for AI* (Now, 2009).

69. Montufar, G. & Morton, J. When does a mixture of products contain a product of mixtures? *J. Discrete Math.* **29**, 321–347 (2014).

70. Montufar, G. F., Pascanu, R., Cho, K. & Bengio, Y. On the number of linear regions of deep neural networks. In *Proc. Advances in Neural Information Processing Systems* **27** 2924–2932 (2014).

71. Bengio, Y., Ducharme, R. & Vincent, P. A neural probabilistic language model. In *Proc. Advances in Neural Information Processing Systems* **13** 932–938 (2001). **This paper introduced neural language models, which learn to convert a word symbol into a word vector or word embedding composed of learned semantic features in order to predict the next word in a sequence.**

72. Cho, K. et al. Learning phrase representations using RNN encoder-decoder

for statistical machine translation. In *Proc. Conference on Empirical Methods in Natural Language Processing* 1724–1734 (2014).

73. Schwenk, H. Continuous space language models. *Computer Speech Lang.* **21**, 492–518 (2007).
74. Socher, R., Lin, C.-Y., Manning, C. & Ng, A. Y. Parsing natural scenes and natural language with recursive neural networks. In *Proc. International Conference on Machine Learning* 129–136 (2011).
75. Mikolov, T., Sutskever, I., Chen, K., Corrado, G. & Dean, J. Distributed representations of words and phrases and their compositionality. In *Proc. Advances in Neural Information Processing Systems* 26 3111–3119 (2013).
76. Bahdanau, D., Cho, K. & Bengio, Y. Neural machine translation by jointly learning to align and translate. In *Proc. International Conference on Learning Representations* <http://arxiv.org/abs/1409.0473> (2015).
77. Hochreiter, S. Untersuchungen zu dynamischen neuronalen Netzen [in German] Diploma thesis, T.U. Münich (1991).
78. Bengio, Y., Simard, P. & Frasconi, P. Learning long-term dependencies with gradient descent is difficult. *IEEE Trans. Neural Networks* **5**, 157–166 (1994).
79. Hochreiter, S. & Schmidhuber, J. Long short-term memory. *Neural Comput.* **9**, 1735–1780 (1997).
- This paper introduced LSTM recurrent networks, which have become a crucial ingredient in recent advances with recurrent networks because they are good at learning long-range dependencies.
80. ElHilli, S. & Bengio, Y. Hierarchical recurrent neural networks for long-term dependencies. In *Proc. Advances in Neural Information Processing Systems* 8 <http://papers.nips.cc/paper/1102-hierarchical-recurrent-neural-networks-for-long-term-dependencies> (1995).
81. Sutskever, I. *Training Recurrent Neural Networks*. PhD thesis, Univ. Toronto (2012).
82. Pascanu, R., Mikolov, T. & Bengio, Y. On the difficulty of training recurrent neural networks. In *Proc. 30th International Conference on Machine Learning* 1310–1318 (2013).
83. Sutskever, I., Martens, J. & Hinton, G. E. Generating text with recurrent neural networks. In *Proc. 28th International Conference on Machine Learning* 1017–1024 (2011).
84. Lakoff, G. & Johnson, M. *Metaphors We Live By* (Univ. Chicago Press, 2008).
85. Rogers, T. T. & McClelland, J. L. *Semantic Cognition: A Parallel Distributed Processing Approach* (MIT Press, 2004).
86. Xu, K. et al. Show, attend and tell: Neural image caption generation with visual attention. In *Proc. International Conference on Learning Representations* <http://arxiv.org/abs/1502.03044> (2015).
87. Graves, A., Mohamed, A.-R. & Hinton, G. Speech recognition with deep recurrent neural networks. In *Proc. International Conference on Acoustics, Speech and Signal Processing* 6645–6649 (2013).
88. Graves, A., Wayne, G. & Danihelka, I. Neural Turing machines. <http://arxiv.org/abs/1410.5401> (2014).
89. Weston, J., Chopra, S. & Bordes, A. Memory networks. <http://arxiv.org/abs/1410.3916> (2014).
90. Weston, J., Bordes, A., Chopra, S. & Mikolov, T. Towards AI-complete question answering: a set of prerequisite toy tasks. <http://arxiv.org/abs/1502.05698> (2015).
91. Hinton, G. E., Dayan, P., Frey, B. J. & Neal, R. M. The wake-sleep algorithm for unsupervised neural networks. *Science* **268**, 1558–1161 (1995).
92. Salakhutdinov, R. & Hinton, G. Deep Boltzmann machines. In *Proc. International Conference on Artificial Intelligence and Statistics* 448–455 (2009).
93. Vincent, P., Larochelle, H., Bengio, Y. & Manzagol, P.-A. Extracting and composing robust features with denoising autoencoders. In *Proc. 25th International Conference on Machine Learning* 1096–1103 (2008).
94. Kavukcuoglu, K. et al. Learning convolutional feature hierarchies for visual recognition. In *Proc. Advances in Neural Information Processing Systems* 23 1090–1098 (2010).
95. Gregor, K. & LeCun, Y. Learning fast approximations of sparse coding. In *Proc. International Conference on Machine Learning* 399–406 (2010).
96. Ranzato, M., Mnih, V., Susskind, J. M. & Hinton, G. E. Modeling natural images using gated MRFs. *IEEE Trans. Pattern Anal. Machine Intell.* **35**, 2206–2222 (2013).
97. Bengio, Y., Thibodeau-Laufer, E., Alain, G. & Yosinski, J. Deep generative stochastic networks trainable by backprop. In *Proc. 31st International Conference on Machine Learning* 226–234 (2014).
98. Kingma, D., Rezende, D., Mohamed, S. & Welling, M. Semi-supervised learning with deep generative models. In *Proc. Advances in Neural Information Processing Systems* 27 3581–3589 (2014).
99. Ba, J., Mnih, V. & Kavukcuoglu, K. Multiple object recognition with visual attention. In *Proc. International Conference on Learning Representations* <http://arxiv.org/abs/1412.7755> (2014).
100. Mnih, V. et al. Human-level control through deep reinforcement learning. *Nature* **518**, 529–533 (2015).
101. Bottou, L. From machine learning to machine reasoning. *Mach. Learn.* **94**, 133–149 (2014).
102. Vinyals, O., Toshev, A., Bengio, S. & Erhan, D. Show and tell: a neural image caption generator. In *Proc. International Conference on Machine Learning* <http://arxiv.org/abs/1502.03044> (2014).
103. van der Maaten, L. & Hinton, G. E. Visualizing data using t-SNE. *J. Mach. Learn. Research* **9**, 2579–2605 (2008).

Acknowledgements The authors would like to thank the Natural Sciences and Engineering Research Council of Canada, the Canadian Institute For Advanced Research (CIFAR), the National Science Foundation and Office of Naval Research for support. Y.L. and Y.B. are CIFAR fellows.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of this paper at go.nature.com/7cjbaa. Correspondence should be addressed to Y.L. (yann@cs.nyu.edu).

Reinforcement learning improves behaviour from evaluative feedback

Michael L. Littman¹

Reinforcement learning is a branch of machine learning concerned with using experience gained through interacting with the world and evaluative feedback to improve a system's ability to make behavioural decisions. It has been called the artificial intelligence problem in a microcosm because learning algorithms must act autonomously to perform well and achieve their goals. Partly driven by the increasing availability of rich data, recent years have seen exciting advances in the theory and practice of reinforcement learning, including developments in fundamental technical areas such as generalization, planning, exploration and empirical methodology, leading to increasing applicability to real-life problems.

Reinforcement-learning algorithms^{1,2} are inspired by our understanding of decision making in humans and other animals in which learning is supervised through the use of reward signals in response to the observed outcomes of actions. As our understanding of this class of problems improves, so does our ability to bring it to bear in practical settings. Reinforcement learning is having a considerable impact on nuts-and-bolts questions such as how to create more effective personalized Web experiences, as well as esoteric questions such as how to design better computerized players for the traditional board game Go or 1980s video games. Reinforcement learning is also providing a valuable conceptual framework for work in psychology, cognitive science, behavioural economics and neuroscience that seeks to explain the process of decision making in the natural world.

One way to think about machine learning is as a set of techniques to try to answer the question: when I am in situation x , what response should I choose? As a concrete example, consider the problem of assigning patrons to tables in a restaurant. Parties of varying sizes arrive in an unknown order and the host allocates each party to a table. The host maps the current situation x (size of the latest party and information about which tables are occupied and for how long) to a decision (which table to assign to the party), trying to satisfy a set of competing goals such as minimizing the waiting time for a table, balancing the load among the various servers, and ensuring that the members of a party can sit together. Similar allocation challenges come up in games such as Tetris and computational problems such as data-centre task allocation. Viewed more broadly, this framework fits any problem in which a sequence of decisions needs to be made to maximize a scoring function over uncertain outcomes.

The kinds of approaches needed to learn good behaviour for the patron-assignment problem depend on what kinds of feedback information are available to the decision maker while learning (Fig. 1).

Exhaustive versus sampled feedback is concerned with the coverage of the training examples. A learner given exhaustive feedback is exposed to all possible situations. Sampled feedback is weaker, in that the learner is only provided with experience of a subset of situations. The central problem in classic supervised learning is generalizing from sampled examples.

Supervised versus evaluative feedback is concerned with how the learner is informed of right and wrong answers. A requirement for applying supervised learning methods is the availability of examples with known optimal decisions. In the patron-assignment problem, a

host-in-training could work as an apprentice to a much more experienced supervisor to learn how to handle a range of situations. If the apprentice can only learn from supervised feedback, however, she would have no opportunities to improve after the apprenticeship ends.

By contrast, evaluative feedback provides the learner with an assessment of the effectiveness of the decisions that she made; no information is available on the appropriateness of alternatives. For example, a host might learn about the ability of a server to handle unruly patrons by trial and error: when the host makes an assignment of a difficult customer, it is possible to tell whether things went smoothly with the selected server, but no direct information is available as to whether one of the other servers might have been a better choice. The central problem in the field of reinforcement learning is addressing the challenge of evaluative feedback.

One-shot versus sequential feedback is concerned with the relative timing of learning signals. Evaluative feedback can be subdivided into whether it is provided directly for each decision or whether it has longer-term impacts that are evaluated over a sequence of decisions. For example, if a host needs to seat a party of 12 and there are no large tables available, some past decision to seat a small party at a big table might be to blame. Reinforcement learners must solve this temporal credit assignment problem to be able to derive good behaviour in the face of this weak sequential feedback.

From the beginning, reinforcement-learning methods have contended with all three forms of weak feedback simultaneously — sampled, evaluative and sequential feedback. As such, the problem is considerably harder than that of supervised learning. However, methods that can learn from weak sources of feedback are more generally applicable and can be mapped to a variety of naturally occurring problems, as suggested by the patron-assignment problem.

The following sections describe recent advances in several sub-areas of reinforcement learning that are expanding its power and applicability.

Bandit problems

The k -armed bandit problem concerns learning to make decisions from one-shot, exhaustive evaluative feedback³. It is both harder (evaluative versus supervised feedback) and easier (exhaustive versus sampled feedback) than supervised learning. Initially, the problem was motivated by decision making in medical trials — prescribing a treatment to a test subject only reveals the outcome of that experiment and not the outcome of competing treatments. Decision making in this setting involves

¹Department of Computer Science, Brown University, Providence, Rhode Island 02912, USA.

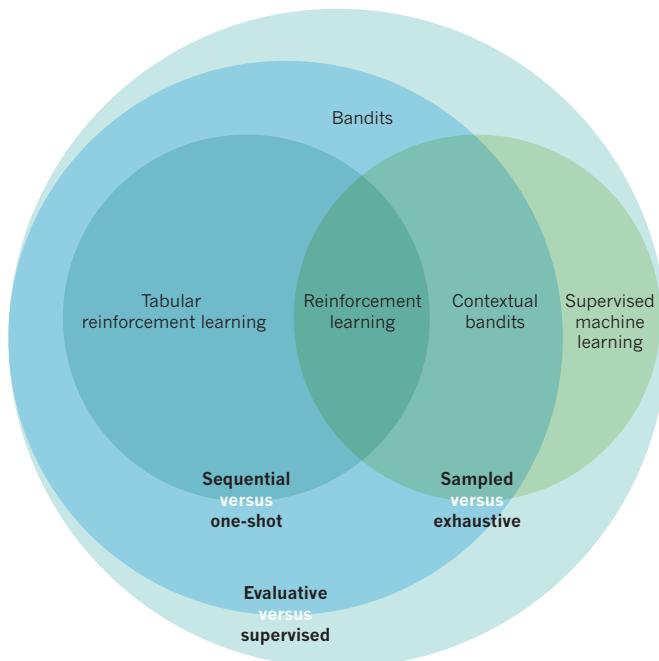


Figure 1 | Decisions of machine-learning feedback. Three kinds of feedback (sequential versus one-shot, sampled versus exhaustive and evaluative versus supervised) in machine learning and examples of learning problems (bandits, tabular reinforcement learning, reinforcement learning, contextual bandits and supervised machine learning) that result from their combination.

an inherent exploration–exploitation trade-off as the learner is trying to perform well, but also needs to test out courses of action that could possibly be better than the current best guess.

With increasing attention being paid to personalized medicine⁴, bandit algorithms could become popular again in the health-care setting. However, the main driver of recent interest in these algorithms is the problem of online content delivery. A company providing web pages to a large user base has many decisions it can make concerning the organization of the information it is providing. Should the page include a large image or a smaller image with space for a brief textual description? A simple formulation of this decision is as a two-armed bandit problem. Each time the system needs to present a page to the user, it chooses between layout A and layout B. It can then estimate the appropriateness of its decision by whether the user clicks on the presented link and, if so, how long the user spends reading and exploring the resulting pages. This information can then be used to inform future layout decisions.

Good algorithms for this problem have been known for more than 50 years. However, increased attention recently has resulted in new analyses and a deeper understanding of which algorithms are most effective in which situations. One popular approach is the upper confidence bound (UCB) algorithm⁵. The idea of the algorithm is quite simple — on the basis of the number of times each alternative has been tested and the outcomes of these tests, the algorithm estimates the likelihood that a link will be selected (its ‘click-through rate’) for layout A and B separately. Statistical confidence bounds of the form ‘with 95% probability, A has a click-through rate of 5%, plus or minus 10%’ are computed. The upper confidence bound is the sum of the estimated click-through rate and the confidence bound; 15%, in this case. There are two reasons a condition might have a higher upper confidence bound: it has a higher estimated click-through rate (because it seems to be a good option — exploitation is warranted) or it has a wider confidence range (because it has not received as many tests — exploration is warranted). The UCB heuristic, then, is to always choose the option with the highest upper confidence bound resulting in either high reward or valuable data for learning. UCB is related to an earlier algorithm called interval estimation⁶, but it is defined so it achieves excellent guarantees on its long-run

regret — roughly speaking, it minimizes the number of clicks lost owing to showing the wrong version of the page.

The UCB algorithm, as described, pays no attention to context. Its choices compare well with the best selection averaged over the entire population of users. Modern web pages have additional information that can be brought to bear, however. The situation at display time, x , includes facts that the system has about the user — recent pages visited, possibly demographic information and perhaps even a record of the user’s recent purchases. The system can use this information to create the most valuable, interesting or engaging page possible. For example, a news site might select a business or sports headline tailored to the current user’s preferences. We, thus, need a way to make good decisions in spite of both evaluative feedback and sampled feedback, a setting that combines aspects of k -armed bandits and supervised learning, now known as contextual bandits. UCB has been extended for use in this context⁷. Indeed, contextual bandit algorithms are increasingly being used by companies with a major Web presence.

A modern approach to bandit problems with quite deep roots is a method now known as Thompson sampling⁸. It bears some similarity to a phenomenon known in psychology as probability matching⁹; as a bandit algorithm, the idea is to use the observations obtained thus far to maintain a posterior distribution over the click-through rate of the alternatives. When it is time to make a decision, the algorithm samples a click-through rate for each alternative from its corresponding posterior distribution, then selects the alternative that was assigned the highest value. Or, more simply, it chooses among the alternatives proportionally to the probability that each is best. Thompson sampling behaves somewhat like a noisier version of UCB, in that uncertain alternatives or confident-but-high-scoring alternatives are most likely to be selected. It can be used in similar scenarios as UCB and can also be shown to possess similar guarantees on its performance^{10,11}. One highly desirable property, however, is that it provides a direct way of applying the emerging class of non-parametric Bayesian methods¹² to decision making — if you can model a process probabilistically, you can use that model to make a selection. This makes it very well suited to decision making in contextual bandit problems.

Temporal difference learning

Learning from sequential feedback, also known as the temporal credit assignment problem, is an essential challenge faced by decision makers whose choices have both immediate and indirect effects. The essence of the idea can be illustrated through noughts and crosses. Imagine the computer is playing a nought against a strong opponent playing a cross. In a series of boards (Fig. 2), nought makes two moves (B and D) and ultimately loses. Which move should be blamed for the loss? A natural assumption is that both moves participated in the loss and therefore both are equally responsible, or that the move that was closer in time to the loss (D) has more responsibility. In this case, however, it was one of nought’s earlier moves (B) that set the stage for its defeat.

The concept of temporal difference learning¹³ provides an algorithmic way to make predictions about the long-term implications of selections. Over a series of games, the learner can estimate whether she is more likely to win or lose from each of the boards it encounters. We can write $V(A)=0$ to represent the fact that nought is likely to tie from board A, and $V(B)=V(C)=V(D)=V(E)=-1$ to represent the fact that nought should be expected to lose from the other boards. The temporal difference error between two consecutive boards x and x' is $V(x') - V(x)$. Here, the temporal difference error is 0 everywhere except for the transition from A to B — it is the decision that changed nought’s situation from a tie to a loss that should be held responsible for the loss.

Leveraging its utility as a marker of successful and unsuccessful decisions, the temporal difference error can be used as a learning signal for improving the value estimates. In particular, an ideal predictor of value will have the property that the long-term prediction for a situation x should be the same as the expected immediate outcome plus the prediction for the resulting situation x' : $V(x) \approx E_x[r(x') + V(x')]$. Here, $r(x')$

represents the feedback received from the environment for the transition to x' . Classic temporal difference methods set about minimizing the difference between these quantities.

Reinforcement learning in the face of evaluative, sampled and sequential feedback requires combining methods for temporal credit assignment with methods for generalization. Concretely, if the V values are represented with a neural network or similar representation, finding predictions such that $V(x) \approx E_x[r(x') + V(x')]$ can lead to instability and divergence in the learning process^{14,15}.

Recent work¹⁶ has modified the goal of learning slightly by incorporating the generalization process into the learning objective itself. Specifically, consider the goal of seeking V values such that $V(x) \approx \Pi E_x[r(x') + V(x')]$ where Π is the projection of the values resulting from their representation by the generalization method. When feedback is exhaustive and no generalization is used, Π is the identity function and this goal is no different from classic temporal difference learning. However, when linear functions are used to represent values, this modified goal can be used to create provably convergent learning algorithms¹⁷.

This insight was rapidly generalized to non-linear function approximators that are smooth (locally linear)¹⁸, to the control setting¹⁹, and to the use of eligibility traces in the learning process²⁰. This line of work holds a great deal of promise for creating reliable methods for learning effective behaviour from very weak feedback.

Planning

A closely related problem to that of temporal credit assignment is planning. The essence of both of these problems is that when actions have both immediate and situation-altering effects, decisions need to

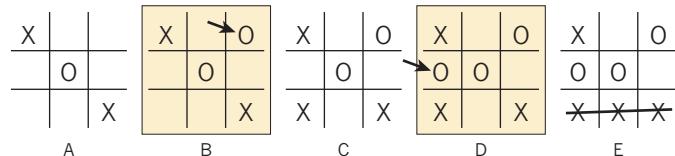


Figure 2 | A series of boards in noughts and crosses leading to a loss for nought. Against a strong opponent, cross (moves A, C and E), nought makes two moves (B and D) and ultimately loses. In this example, nought's earlier move (B) set the stage for the defeat.

be sensitive to both immediate and long-term outcomes. Temporal difference learning provides one mechanism for factoring long-term outcomes into the decision process. Planning is another in that it considers the implications of taking entire sequences of actions and explicitly makes a trade-off among them. Planning has long been studied in artificial intelligence²¹, but research in reinforcement learning has brought about a number of powerful innovations.

A particularly challenging kind of planning that is relevant in the reinforcement-learning setting is planning under uncertainty. Here, outcomes of decisions are modelled as being drawn from a probability distribution over alternatives conditioned on the learner's decision (Box 1). A familiar example of this kind of planning is move selection in board games. If a computer program is playing against a human, it needs to make moves that are likely to lead to a win despite not knowing precisely which moves its opponent will take. For a game such as noughts and crosses, a computer program can exhaustively enumerate all the reachable boards — the so-called game tree — and calculate

BOX 1

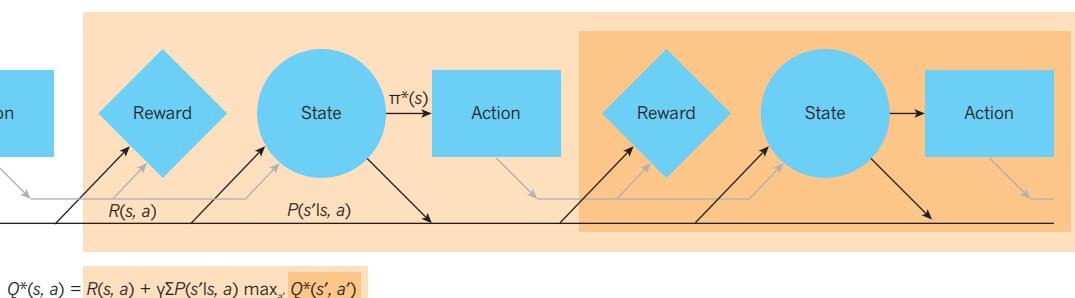
Markov decision processes specify setting and tasks

Reinforcement learning is characterized as an interaction between a learner and an environment that provides evaluative feedback. The environment is often conceptualized as a Markov decision process — a formal mathematical model that dates back to the 1950s (refs 55, 56). A Markov decision process is defined by a set of states S (situations in which a decision can be made), and actions A (the decisions or interventions the decision maker can select). These quantities can be taken to be finite, but continuous state and action spaces are often valuable for capturing interactions in important reinforcement-learning applications such as robotic control⁵⁷. A transition function $P(s'|s, a)$ defines the probability of the state changing from s to s' under the influence of action a . It specifies the ‘physics’ or ‘dynamics’ of the environment.

The decision maker’s task is defined by a reward function $R(s, a)$ and discount factor $\gamma \in [0, 1]$. Rewards are delivered to the decision maker with each transition and maximizing the cumulative discounted expected reward is its objective. Specifically, the decision maker seeks a behaviour π^* mapping states to actions generating a sequence of rewards $r_0, r_1, r_2, r_3, \dots$ such that $E_{r_0, r_1, \dots} [r_0 + \gamma r_1 + \gamma^2 r_2 + \gamma^3 r_3 + \dots]$ is as

large as possible. The relation between the environmental interaction (state, action, reward, state, action, reward, ...) and the cumulative discounted expected reward is captured by the Bellman equation (Figure) for the optimal state-action value function Q^* . The solution to the Bellman equation can be used to define optimal behaviour by $\pi^*(s) = \arg \max_a Q^*(s, a)$. The cumulative discounted expected reward for the policy that takes action a from state s and then behaving optimally thereafter is the immediate reward received plus the discounted expected value of the cumulative discounted expected reward from the resulting state s' given that the best action is chosen.

Planning methods use knowledge of P and R to compute a good policy π . Reinforcement-learning methods have access to P and R only through their ongoing interactions with the environment and must learn good behaviour. Note that the Markov decision process model captures both sequential feedback and the more specific one-shot feedback (when $P(s'|s, a)$ is independent of s and a). It also captures both exhaustive feedback and the more general sampled feedback (when states are represented by features that allow $R(s, a)$ to be represented compactly).



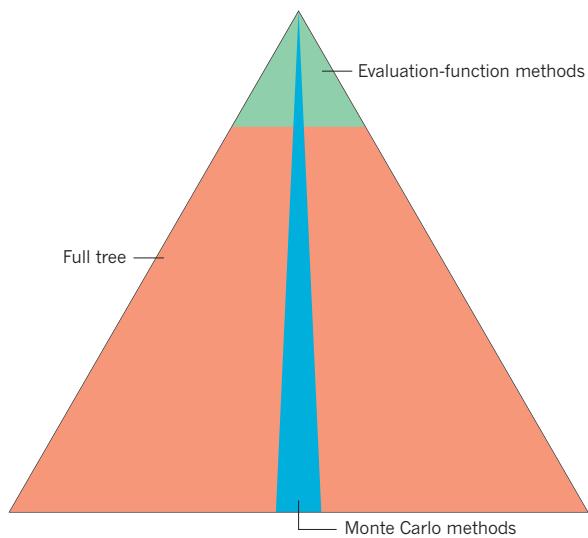


Figure 3 | Schematic comparison of evaluation-function and Monte Carlo methods for planning. The top of the tree represents the current state for which a decision is needed. The red area represents the entire search tree — too large to be examined. The green area represents the nodes of the tree visited by evaluation-function methods — they expand to a fixed depth and then use an evaluation function to approximate the value received at the end of the game. The blue area represents the nodes of the tree visited by Monte Carlo methods — they consider a subset of nodes that ignore some actions (making it narrow) but continue to the leaves of the search tree (making it deep). Some decision problems are much more amenable to this type of approximation.

whether its moves will lead to a win, a loss or a tie from each. In a game such as chess, however, the size of the game tree is astronomical and the best an algorithm can hope for is to assemble a representative sample of boards to consider. The classic approach for dealing with this problem is to use an evaluation function. Programs search as deeply in the game tree as they can, but if they cannot reach the end of the game where the outcome is known, a heuristic evaluation function is used to estimate how the game will end. Even a moderately good guess can lead to excellent game play in this context — evaluation-function methods have been used in chess from the earliest days of artificial intelligence to the decision making in the Deep Blue system that beat the World Chess Champion Garry Kasparov²².

A powerful variation of this idea is to use machine learning to improve the evaluation function over time. Here, temporal difference

learning can be used in simulated games to provide opportunities to learn an accurate evaluation function represented as, for example, a neural network. Celebrated examples of the use of temporal difference methods combined with generalization were in the domains of checkers²³ and backgammon²⁴. A similar approach was responsible for Daily Double wagering in the Watson system for playing the US television game show Jeopardy!²⁵.

In spite of their impressive, indeed superhuman, performance in some games, evaluation-function methods have not been universally successful. A notable example is the game of Go, in which good evaluation functions have been hard to come by, resulting in subpar game play. A relatively recent breakthrough was the introduction of Monte Carlo tree search (MCTS) methods that swap an estimate of the evaluation function for an estimate of a good policy. Whereas evaluation-function methods try all possible actions out to a fixed depth and estimate the value of the resulting boards, Monte Carlo methods search all the way to the end of the game where evaluations are exact, but use a heuristic rule to suggest which subset of actions are worth considering (Fig. 3). Use of the upper confidence intervals in trees (UCT) algorithm²⁶ resulted in a quantum leap in performance in Go²⁷. It leverages the UCB algorithm to intelligently allocate search resources at each point in the tree.

Model-based reinforcement learning

To plan, a learning algorithm needs to be able to predict the immediate outcomes of its actions. Since it can try an action and observe the effects right away, this part of the problem can be addressed using supervised learning methods. In model-based reinforcement-learning approaches, experiences are used to estimate a transition model, mapping from situations and decisions to resulting situations, and then a planning approach is used to make decisions that maximize predicted long-term outcomes.

Model-based reinforcement-learning methods and model-free methods such as temporal differences strike a different balance between computational expense and experience efficiency — to a first approximation, model-free methods are inexpensive, but slow to learn; whereas model-based methods are computationally intensive, but squeeze the most out of each experience (Box 2). In applications such as learning to manage memory resources in a central processing unit, decisions need to be made at the speed of a computer's clock and data are abundant, making temporal difference learning an excellent fit²⁸. In applications such as a robot helicopter learning acrobatic tricks, experience is slow to collect and there is ample time offline to compute decision policies, making a model-based approach appropriate²⁹. There are also methods that draw on both styles of learning³⁰ to try to get the complementary advantages.

Model-based algorithms can make use of insights from supervised learning to deal with sampled feedback when learning their transition

BOX 2

Two approaches to tabular reinforcement learning

Tabular reinforcement-learning problems are those in which the state-action space can be explored exhaustively. Although challenging, there are a number of relatively simple algorithms that can be used effectively in this setting. Q-learning⁵⁸ is a model-free algorithm that uses transition experience of the form $\langle s, a, r, s' \rangle$ (from state s , action a resulted in reward r and next state s') to improve an estimate \hat{Q} of the optimal state-action value function Q^* . In particular, since $r + \gamma \max_a Q^*(s', a)$ is an unbiased estimator of $Q^*(s, a)$ as defined by the Bellman equation and \hat{Q} is an estimate of Q^* , the estimate can be updated as:

$$\hat{Q}(s, a) \leftarrow \hat{Q}(s, a) + \alpha(r + \gamma \max_{a'} \hat{Q}(s', a') - \hat{Q}(s, a)),$$

where α is a learning-rate parameter that controls how new estimates are iteratively blended together over time. Indeed,

if the learning rate is decayed at an appropriate rate,

$$\hat{Q}(s, a) \rightarrow Q^*(s, a)$$

for all state-action pairs⁵⁹.

A simple model-based approach to tabular reinforcement learning uses the observed transition experience to estimate $R(s, a)$ (by averaging all rewards received from state s and action a) and $P(s'|s, a)$ (by the fraction of transitions from state s and action a that transition to state s'). Solving the Bellman equation — through planning algorithms such as value iteration, policy iteration, UCT (upper confidence intervals in trees) or linear programming methods — using these estimates leads to an approximation \hat{Q} of the state-action value function that converges to Q^* in the limit of infinite data.

BOX 3

Self-aware learning in reinforcement learning

The KWIK ('knows what it knows') learning setting³³ provides a powerful set of algorithms for rapidly learning accurate predictions and confidently distinguishing learned from not-yet-learned instances. In the context of reinforcement-learning algorithms, this ability to tag predictions as uncertain can be exploited to drive exploration by substituting optimistic estimates for those that remain unknown.

A growing set of KWIK-learnable function classes has been identified. As a simple example, consider trying to learn a function of the form $f(x) = \text{is } x \text{ divisible by } k?$ for some unknown natural number k . The Table shows what a KWIK learning algorithm for this problem would predict given training examples $f(6) = \text{yes}$ and $f(20) = \text{no}$. It would know $f(6) = \text{yes}$ and $f(20) = \text{no}$ from these examples. It would be able to reason that the unknown $k \in \{2, 3, 6\}$ from these examples and therefore that $f(25) = \text{no}$. By contrast, $f(15) = \text{I don't know}$ because it would be 'yes' if $k=3$ and 'no' if $k=6$. Despite this lingering ambiguity, it could conclude that $f(24) = \text{yes}$, since that would be the answer for any of the remaining possible values of k . This last element lies at the heart of all KWIK-learning algorithms — not 'wasting' its I don't know answers on examples whose outputs can be inferred from the data.

For noisy function classes, a KWIK learner must be able to predict the probability of various outputs given an input. This property is important when learning Markov decision process transition functions. Some challenging function classes for which KWIK learning algorithms are known include noisy linear regression, and noisy union. Noisy union, also called the 'adaptive k meteorologist' problem⁵⁰, lets the learner choose a function from a finite class when there is a non-zero probability that training examples are corrupted.

Table 1 | A training set ($f(6) = \text{yes}$, $f(20) = \text{no}$) for a function of the form ' $f(x) = \text{is } x \text{ divisible by } k?$ ' for some unknown natural number k , and the predictions made by a valid KWIK algorithm.

Input x	Output $f(x)$	Prediction for $f(x)$
6	Yes	Yes
15		I don't know
20	No	No
24		Yes
25		No

models and can tackle the problem of sequential feedback by planning actions to achieve their goals. The issue of evaluative feedback — the fact that the system can only learn about transitions it has actually experienced — makes it very important that learners balance exploration and exploitation. Several methods are known for achieving guaranteed efficiency in model-based reinforcement learning^{31,32} by explicitly modelling the certainty of predictions and, much like UCB, taking an optimistic view of the uncertainty. Concretely representing the uncertainty in learning leads to the framework 'knows what it knows'³³ (Box 3), which has been shown to work well with planning algorithms in model-based reinforcement-learning settings.

Empirical methods

An important turning point in the development of practical supervised learning methods was a shift from artificial learning problems to learning problems grounded in measured data and careful experimental design³⁴. The reinforcement-learning community is showing signs of a similar evolution, which promises to help to solidify the technical gains described in the previous sections. Of particular interest are new evaluation methodologies that go beyond case studies of individual algorithm learning in individual environments to controlled studies of different learning algorithms and their performance across multiple natural domains. In supervised learning, data-set collections such as the University of California, Irvine Machine Learning Repository (<http://archive.ics.uci.edu/ml/>) and the Weka Machine Learning library (<http://www.cs.waikato.ac.nz/ml/weka/>) facilitate this kind of research by providing data for learning problems and a suite of algorithms in a modular form. A challenge for the study of reinforcement learning, however, is that learning takes place through active exploration between the learner and her environment. Environments are almost always represented by simulators, which can be difficult to standardize and share, and are rarely realistic stand-ins for the actual environment of interest.

One noteworthy exception is the Arcade Learning Environment³⁵, which provides a reinforcement-leaning-friendly interface for an eclectic collection of more than 50 Atari 2600 video games. Although video games are by no means the only domain of interest, the diversity of game types and the unity of the interface makes the Arcade Learning Environment an exciting platform for the empirical study of machine reinforcement learning. An extremely impressive example of

research enabled by this repository is a recent study of deep learning in the reinforcement-learning setting. Researchers combined a set of existing and novel algorithm elements to create an agent capable of 'human-level' performance³⁶ across the set of games. It remains to be seen whether the algorithmic approaches highlighted in this work become standard weapons in the reinforcement-learning arsenal, but the remarkable success on these challenging learning problems has attracted a great deal of attention and is likely to generate advances in the theory and practice of reinforcement learning.

A perspective on future impacts

There are several current directions that promise considerable impacts in the years to come.

Offline evaluation

One challenge to designing practical, reliable learning systems for real-life reinforcement-learning problems is that the learning problem is carried out online. That is, in contrast to supervised learning systems that work with a training set that contains measured, but frozen, data, reinforcement learning is defined with regard to a continual interaction between the learner and an environment.

To overcome this difficulty, several efforts are underway to collect static data sets in a way that allows them to be used for evaluating reinforcement-learning algorithms. In the biostatistics community, a new type of clinical trial³⁷ called a sequential, multiple assignment randomized trial (SMART) has been developed that makes it possible to gather treatment responses in a way that supports the synthesis of adaptive treatment plans. That is, the trial assesses the effectiveness of treatments expressed as flow charts in which an intervention is applied and, depending on the patient's response, the plan follows up with different possible interventions. A typical SMART can evaluate a flow chart containing eight different possible treatment plans using half the participants that would otherwise be needed.

SMARTs evaluate flow-chart-like treatment plans, but not complete learning algorithms. Learning algorithms can express extremely complex contingent relationships between the data acquired during learning and the decisions made in response to it. Consider a simple two-armed bandit problem over a set of ten steps. The tenth decision in the sequence is conditioned on the outcomes of all the previous decisions. Thus, for each of the $2^9 = 512$ possible contexts, our bandit algorithm needs to make

a good choice. Longer learning periods and contextual decisions make the number of possibilities to consider astronomically high. Despite this daunting combinatorial explosion, offline evaluation procedures for contextual bandit problems have been proposed³⁸. The key idea is to collect a static (but enormous) set of contextual bandit decisions by running a uniform random selection algorithm on a real problem (a news article recommendation, for example). Then, to test a bandit algorithm using this collection, present the algorithm with the same series of contexts. Whenever the selection made by the algorithm does not match the selection made during data collection, the algorithm is forced to forget what it saw. Thus, as far as the algorithm ‘knows’, all of its choices during learning matched the choices made during data collection. The end result is an unbiased evaluation of a dynamic bandit algorithm using static pre-recorded data.

One reason that this approach works for contextual bandit problems is that — apart from the state of the learning algorithm — there are no dependencies from one round of decision making to the next. So, skipping many rounds because the algorithm and the data collection do not match does not change the fundamental character of the learning problem. In the presence of sequential feedback, however, this trick no longer applies. Proposals have been made for creating reinforcement-learning data repositories³⁹ and for evaluating policies using observational data sets⁴⁰, but no solid solution for evaluating general learning algorithms has yet emerged.

Finding appropriate reward functions

Reinforcement-learning systems strive to identify behaviour that maximizes expected total reward. There is a sense, therefore, that they are ‘programmable’ through their reward functions — the mappings from situation to score. A helpful analogy to consider is between learning and program interpretation in digital computers: a program (reward function) directs an interpreter (learning algorithm) to process inputs (environments) into desirable outputs (behaviour). The vast majority of research in reinforcement learning has been into the development of effective interpreters with little attention paid to how we should be programming them (providing reward functions).

There are a few approaches that have been explored for producing reward functions that induce a target behaviour given some conception of that behaviour. As a simple example, if one reward function for the target behaviour is known, the space of behaviour-preserving transformations to this reward function is well understood⁴¹; other, essentially equivalent, reward functions can be created. If a human teacher is available who can provide evaluative feedback, modifications to intended target reinforcement-learning algorithms can be used to search for the target behaviour^{42–44}. The problem of inverse reinforcement learning^{45,46} addresses the challenge of creating appropriate reward functions given the availability of behavioural traces from an expert executing the target behaviour. It is called ‘inverse’ reinforcement learning because the learner is given behaviour and needs to generate a reward function instead of the other way around. These methods infer a reward function that the demonstrated behaviour optimizes. One can turn to evolutionary optimization⁴⁷ to generate reward functions if there is an effective way to evaluate the appropriateness of the resulting behaviour. One advantage of the evolutionary approach is that, among the many possible reward functions that generate good behaviour, it can identify ones that provide helpful but not too distracting hints that can speed up the learning process.

Looking forward, new techniques for specifying complex behaviour and translations of these specifications into appropriate reward functions are essential. Existing reward-function specifications lack the fundamental ideas that enable the design of today’s massive software systems, such as abstraction, modularity and encapsulation. Analogues of these ideas could greatly extend the practical utility of reinforcement-learning systems.

Reinforcement learning as a cognitive model

The term reinforcement learning itself originated in the animal-learning community. It had long been observed that certain stimuli,

if present after an animal selects a particular behavioural action, cause that behavioural action to become more frequent in the future. These stimuli became known as reinforcers because of their role in strengthening behaviour. The pioneers of artificial intelligence were aware of this idea^{48,49} and used the term to describe computational processes that would similarly increase the intensity or likelihood of a response. Indeed, according to the software tool Google Ngrams (which charts the frequency of phrases in its database of published books; <https://books.google.com/ngrams>), the appearance of ‘reinforcement learning’ reached a local maximum in the mid-1960s, as it was in use by researchers studying the theory of animal learning, education and cybernetics.

The term lost ground through the 1970s and early 1980s, but was picked up again by computer scientists in the mid-to-late 1980s^{13,50} and became a central area of study in machine learning. A subtle but important difference was introduced during this second wave. The term ‘reinforcement’ came to refer to a generic reward signal that needs to be predicted over time instead of the activity of strengthening a behavioural response. Classic learning problems could still be discussed in this way, but the door opened to a wide variety of other powerful algorithmic approaches, such as temporal difference learning, for which behavioural strengthening was not a direct design goal.

As already discussed, the predictive view proved to be powerful and productive. Recent work in neuroscience and cognitive psychology has leveraged this revised understanding of reinforcement learning to create new insights into biological learning. Of particular note is the discovery of a temporal difference signal encoded in the firing of dopamine neurons⁵¹, which has created a bridge between computer scientists and cognitive neuroscientists centred on reinforcement-learning algorithms^{52,53}. In addition to temporal difference methods, cognitive scientists have found the distinction between model-based (goal-directed) and model-free (habitual) learning to be a fertile one. For example, some aspects of how people make moral judgments⁵⁴ can be fruitfully described by associating action-based preferences (do not do A because it is wrong) with model-free learning and outcome-based preferences (do not do A because it will lead to a bad situation) with model-based learning.

In October 2013, the first Multi-disciplinary Conference on Reinforcement Learning and Decision Making was held at Princeton University in New Jersey. In addition to computer scientists, psychologists and neuroscientists, the invited speakers included engineers, roboticists, a zoologist, a psychiatrist and a behavioural game theorist. The talks were notable both for the breadth of topics and the convergence of the underlying conceptual framework discussed in this Review. As our formal and practical understanding of learning from evaluative feedback expands, so does the community of scientists and engineers who can contribute to and benefit from these ideas. ■

Received 11 January; accepted 28 April 2015.

1. Sutton, R. S. & Barto, A. G. *Reinforcement Learning: An Introduction* (MIT Press, 1998). **This book is the definitive reference on computational reinforcement learning.**
2. Kaelbling, L. P., Littman, M. L. & Moore, A. W. Reinforcement learning: a survey. *J. Artif. Intell. Res.* **4**, 237–285 (1996).
3. Berry, D. A. & Fristedt, B. *Bandit Problems: Sequential Allocation of Experiments* (Chapman and Hall, 1985).
4. Shrager, J. & Tenenbaum, J. M. Rapid learning for precision oncology. *Nature Rev. Clin. Onco.* **11**, 109–118 (2014).
5. Auer, P., Cesa-Bianchi, N. & Fischer, P. Finite-time analysis of the multi-armed bandit problem. *Mach. Learn.* **47**, 235–256 (2002).
6. Kaelbling, L. P. *Learning in Embedded Systems* (MIT Press, 1993).
7. Li, L., Chu, W., Langford, J. & Schapire, R. E. A contextual-bandit approach to personalized news article recommendation. In *Proc. 19th International World Wide Web Conference* 661–670 (2010).
8. Thompson, W. R. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* **25**, 285–294 (1933).
9. West, R. F. & Stanovich, K. E. Is probability matching smart? Associations between probabilistic choices and cognitive ability. *Mem. Cognit.* **31**, 243–251 (2003).
10. May, B. C., Korda, N., Lee, A. & Leslie, D. S. Optimistic Bayesian sampling in contextual-bandit problems. *J. Mach. Learn. Res.* **13**, 2069–2106 (2012).

11. Bubeck, S. & Liu, C.-Y. Prior-free and prior-dependent regret bounds for Thompson sampling. In *Proc. Advances in Neural Information Processing Systems* 638–646 (2013).

12. Gershman, S. & Blei, D. A tutorial on Bayesian nonparametric models. *J. Math. Psychol.* **56**, 1–12 (2012).

13. Sutton, R. S. Learning to predict by the method of temporal differences. *Mach. Learn.* **3**, 9–44 (1988).

14. Boyan, J. A. & Moore, A. W. Generalization in reinforcement learning: safely approximating the value function. In *Proc. Advances in Neural Information Processing Systems* 369–376 (1995).

15. Baird, L. Residual algorithms: reinforcement learning with function approximation. In *Proc. 12th International Conference on Machine Learning* (eds Priditidis, A. & Russell, S.) 30–37 (Morgan Kaufmann, 1995).

16. Sutton, R. S. *et al.* Fast gradient-descent methods for temporal-difference learning with linear function approximation. In *Proc. 26th Annual International Conference on Machine Learning* 993–1000 (2009).

17. Sutton, R. S., Maei, H. R. & Szepesvári, C. A convergent O(n) temporal-difference algorithm for off-policy learning with linear function approximation. In *Proc. Advances in Neural Information Processing Systems* 1609–1616 (2009).

18. Maei, H. R. *et al.* Convergent temporal-difference learning with arbitrary smooth function approximation. In *Proc. Advances in Neural Information Processing Systems* 1204–1212 (2009).

19. Maei, H. R., Szepesvári, C., Bhatnagar, S. & Sutton, R. S. Toward off-policy learning control with function approximation. In *Proc. 27th International Conference on Machine Learning* 719–726 (2010).

20. van Hasselt, H., Mahmood, A. R. & Sutton, R. S. Off-policy TD(λ) with a true online equivalence. In *Proc. 30th Conference on Uncertainty in Artificial Intelligence* 324 (2014).

21. Russell, S. J. & Norvig, P. *Artificial Intelligence: A Modern Approach* (Prentice-Hall, 1994).

22. Campbell, M., Hoane, A. J. & Hsu, F. H. Deep blue. *Artif. Intell.* **134**, 57–83 (2002).

23. Samuel, A. L. Some studies in machine learning using the game of checkers. *IBM J. Res. Develop.* **3**, 211–229 (1959).

24. Tesauro, G. TD-Gammon, a self-teaching backgammon program, achieves master-level play. *Neural Comput.* **6**, 215–219 (1994).

This article describes the first reinforcement-learning system to solve a truly non-trivial task.

25. Tesauro, G., Gondek, D., Lenchner, J., Fan, J. & Prager, J. M. Simulation, learning, and optimization techniques in Watson's game strategies. *IBM J. Res. Develop.* **56**, 1–11 (2012).

26. Kocsis, L. & Szepesvári, C. Bandit based Monte-Carlo planning. In *Proc. 17th European Conference on Machine Learning* 282–293 (2006).

This article introduces UCT, the decision-making algorithm that revolutionized gameplay in Go.

27. Gelly, S. *et al.* The grand challenge of computer Go: Monte Carlo tree search and extensions. *Communications of the ACM* **55**, 106–113 (2012).

28. Ipek, E., Mutlu, O., Martínez, J. F. & Caruana, R. Self-optimizing memory controllers: a reinforcement learning approach. In *Proc. 35th International Symposium on Computer Architecture* 39–50 (2008).

29. Ng, A. Y., Kim, H. J., Jordan, M. I. & Sastry, S. Autonomous helicopter flight via reinforcement learning. In *Proc. Advances in Neural Information Processing Systems* <http://papers.nips.cc/paper/2455-autonomous-helicopter-flight-via-reinforcement-learning> (2003).

30. Sutton, R. S. Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In *Proc. 7th International Conference on Machine Learning* 216–224 (Morgan Kaufmann, 1990).

31. Kearns, M. J. & Singh, S. P. Near-optimal reinforcement learning in polynomial time. *Mach. Learn.* **49**, 209–232 (2002).

This article provides the first algorithm and analysis that shows that reinforcement-learning tasks can be solved approximately optimally with a relatively small amount of experience.

32. Brafman, R. I. & Tennenholtz, M. R-MAX — a general polynomial time algorithm for near-optimal reinforcement learning. *J. Mach. Learn. Res.* **3**, 213–231 (2002).

33. Li, L., Littman, M. L., Walsh, T. J. & Strehl, A. L. Knows what it knows: a framework for self-aware learning. *Mach. Learn.* **82**, 399–443 (2011).

34. Langley, P. Machine learning as an experimental science. *Mach. Learn.* **3**, 5–8 (1988).

35. Bellemare, M. G., Naddaf, Y., Veness, J. & Bowling, M. The arcade learning environment: an evaluation platform for general agents. *J. Artif. Intell. Res.* **47**, 253–279 (2013).

36. Mnih, V. *et al.* Human-level control through deep reinforcement learning. *Nature* **518**, 529–533 (2015).

This article describes the application of deep learning in a reinforcement-learning setting to address the challenging task of decision making in an arcade environment.

37. Murphy, S. A. An experimental design for the development of adaptive treatment strategies. *Stat. Med.* **24**, 1455–1481 (2005).

38. Li, L., Chu, W., Langford, J. & Wang, X. Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms. In *Proc. 4th ACM International Conference on Web Search and Data Mining* 297–306 (2011).

39. Nouri, A. *et al.* A novel benchmark methodology and data repository for real-life reinforcement learning. In *Proc. Multidisciplinary Symposium on Reinforcement Learning*, Poster (2009).

40. Marivate, V. N., Chemali, J., Littman, M. & Brunskill, E. Discovering multi-modal characteristics in observational clinical data. In *Proc. Machine Learning for Clinical Data Analysis and Healthcare NIPS Workshop* <http://paul.rutgers.edu/~vukosi/papers/nips2013workshop.pdf> (2013).

41. Ng, A. Y., Harada, D. & Russell, S. Policy invariance under reward transformations: theory and application to reward shaping. In *Proc. 16th International Conference on Machine Learning* 278–287 (1999).

42. Thomaz, A. L. & Breazeal, C. Teachable robots: understanding human teaching behaviour to build more effective robot learners. *Artif. Intell.* **172**, 716–737 (2008).

43. Knox, W. B. & Stone, P. Interactively shaping agents via human reinforcement: The TAMER framework. In *Proc. 5th International Conference on Knowledge Capture* 9–16 (2009).

44. Loftin, R. *et al.* A strategy-aware technique for learning behaviors from discrete human feedback. In *Proc. 28th Association for the Advancement of Artificial Intelligence Conference* <https://www.aaai.org/ocs/index.php/AAAI14/paper/view/8579> (2014).

45. Ng, A. Y. & Russell, S. Algorithms for inverse reinforcement learning. In *Proc. International Conference on Machine Learning* 663–670 (2000).

46. Babes, M., Marivate, V. N., Littman, M. L. & Subramanian, K. Apprenticeship learning about multiple intentions. In *Proc. International Conference on Machine Learning* 897–904 (2011).

47. Singh, S., Lewis, R. L., Barto, A. G. & Sorg, J. Intrinsically motivated reinforcement learning: an evolutionary perspective. *IEEE Trans. Auto. Mental Dev.* **2**, 70–82 (2010).

48. Newell, A. The chess machine: an example of dealing with a complex task by adaptation. In *Proc. Western Joint Computer Conference* 101–108 (1955).

49. Minsky, M. L. Some methods of artificial intelligence and heuristic programming. In *Proc. Symposium on the Mechanization of Thought Processes* 24–27 (1958).

50. Sutton, R. S. & Barto, A. G. Toward a modern theory of adaptive networks: expectation and prediction. *Psychol. Rev.* **88**, 135–170 (1981).

51. Schultz, W., Dayan, P. & Montague, P. R. A neural substrate of prediction and reward. *Science* **275**, 1593–1599 (1997).

52. Dayan, P. & Niv, Y. Reinforcement learning and the brain: the good, the bad and the ugly. *Curr. Opin. Neurobiol.* **18**, 185–196 (2008).

53. Niv, Y. Neuroscience: dopamine ramps up. *Nature* **500**, 533–535 (2013).

54. Cushman, F. Action, outcome, and value a dual-system framework for morality. *Pers. Soc. Psychol. Rev.* **17**, 273–292 (2013).

55. Shapley, L. Stochastic games. *Proc. Natl Acad. Sci. USA* **39**, 1095–1100 (1953).

56. Bellman, R. *Dynamic Programming* (Princeton Univ. Press, 1957).

57. Kober, J., Bagnell, J. A. & Peters, J. Reinforcement learning in robotics: a survey. *Int. J. Rob. Res.* **32**, 1238–1274 (2013).

58. Watkins, C. J. C. H. & Dayan, P. Q-learning. *Mach. Learn.* **8**, 279–292 (1992).

This article introduces the first provably correct approach to reinforcement learning for both prediction and decision making.

59. Jaakkola, T., Jordan, M. I. & Singh, S. P. Convergence of stochastic iterative dynamic programming algorithms. In *Advances in Neural Information Processing Systems* 6, 703–710 (Morgan Kaufmann, 1994).

60. Diuk, C., Li, L. & Leffner, B. R. The adaptive k-meteorologists problem and its application to structure learning and feature selection in reinforcement learning. In *Proc. 26th International Conference on Machine Learning* 32–40 (2009).

Acknowledgements

The author appreciates his discussions with his colleagues that led to this synthesis of current work.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The author declares no competing financial interests. Readers are welcome to comment on the online version of this paper at go.nature.com/cslqwl. Correspondence should be addressed to M.L.L. (mllittman@cs.brown.edu).

Probabilistic machine learning and artificial intelligence

Zoubin Ghahramani¹

How can a machine learn from experience? Probabilistic modelling provides a framework for understanding what learning is, and has therefore emerged as one of the principal theoretical and practical approaches for designing machines that learn from data acquired through experience. The probabilistic framework, which describes how to represent and manipulate uncertainty about models and predictions, has a central role in scientific data analysis, machine learning, robotics, cognitive science and artificial intelligence. This Review provides an introduction to this framework, and discusses some of the state-of-the-art advances in the field, namely, probabilistic programming, Bayesian optimization, data compression and automatic model discovery.

The key idea behind the probabilistic framework to machine learning is that learning can be thought of as inferring plausible models to explain observed data. A machine can use such models to make predictions about future data, and take decisions that are rational given these predictions. Uncertainty plays a fundamental part in all of this. Observed data can be consistent with many models, and therefore which model is appropriate, given the data, is uncertain. Similarly, predictions about future data and the future consequences of actions are uncertain. Probability theory provides a framework for modelling uncertainty.

This Review starts with an introduction to the probabilistic approach to machine learning and Bayesian inference, and then discusses some of the state-of-the-art advances in the field. Many aspects of learning and intelligence crucially depend on the careful probabilistic representation of uncertainty. Probabilistic approaches have only recently become a mainstream approach to artificial intelligence¹, robotics² and machine learning^{3,4}. Even now, there is controversy in these fields about how important it is to fully represent uncertainty. For example, advances using deep neural networks to solve challenging pattern-recognition problems such as speech recognition⁵, image classification^{6,7}, and prediction of words in text⁸, do not overtly represent the uncertainty in the structure or parameters of those neural networks. However, my focus will not be on these types of pattern-recognition problems, characterized by the availability of large amounts of data, but on problems for which uncertainty is really a key ingredient, for example where a decision may depend on the amount of uncertainty.

I highlight five areas of current research at the frontier of probabilistic machine learning, emphasizing areas that are of broad relevance to scientists across many fields: probabilistic programming, which is a general framework for expressing probabilistic models as computer programs and which could have a major impact on scientific modelling; Bayesian optimization, which is an approach to globally optimizing unknown functions; probabilistic data compression; automating the discovery of plausible and interpretable models from data; and hierarchical modelling for learning many related models, for example for personalized medicine or recommendation. Although considerable challenges remain, the coming decade promises substantial advances in artificial intelligence and machine learning based on the probabilistic framework.

Probabilistic modelling and representing uncertainty

At the most basic level, machine learning seeks to develop methods for computers to improve their performance at certain tasks on the basis of

observed data. Typical examples of such tasks might include detecting pedestrians in images taken from an autonomous vehicle, classifying gene-expression patterns from leukaemia patients into subtypes by clinical outcome, or translating English sentences into French. However, as I discuss, the scope of machine-learning tasks is even broader than these pattern classification or mapping tasks, and can include optimization and decision making, compressing data and automatically extracting interpretable models from data.

Data are the key ingredients of all machine-learning systems. But data, even so-called big data, are useless on their own until one extracts knowledge or inferences from them. Almost all machine-learning tasks can be formulated as making inferences about missing or latent data from the observed data — I will variously use the terms inference, prediction or forecasting to refer to this general task. Elaborating the example mentioned, consider classifying people with leukaemia into one of the four main subtypes of this disease on the basis of each person's measured gene-expression patterns. Here, the observed data are pairs of gene-expression patterns and labelled subtypes, and the unobserved or missing data to be inferred are the subtypes for new patients. To make inferences about unobserved data from the observed data, the learning system needs to make some assumptions; taken together these assumptions constitute a model. A model can be very simple and rigid, such as a classic statistical linear regression model, or complex and flexible, such as a large and deep neural network, or even a model with infinitely many parameters. I return to this point in the next section. A model is considered to be well defined if it can make forecasts or predictions about unobserved data having been trained on observed data (otherwise, if the model cannot make predictions it cannot be falsified, in the sense of the philosopher Karl Popper's proposal for evaluating hypotheses, or as the theoretical physicist Wolfgang Pauli said the model is "not even wrong"). For example, in the classification setting, a well-defined model should be able to provide predictions of class labels for new patients. Since any sensible model will be uncertain when predicting unobserved data, uncertainty plays a fundamental part in modelling.

There are many forms of uncertainty in modelling. At the lowest level, model uncertainty is introduced from measurement noise, for example, pixel noise or blur in images. At higher levels, a model may have many parameters, such as the coefficients of a linear regression, and there is uncertainty about which values of these parameters will be good at predicting new data. Finally, at the highest levels, there is

¹Department of Engineering, University of Cambridge, Trumpington Street, Cambridge CB2 1PZ, UK.

often uncertainty about even the general structure of the model: is linear regression or a neural network appropriate, if the latter, how many layers should it have, and so on.

The probabilistic approach to modelling uses probability theory to express all forms of uncertainty⁹. Probability theory is the mathematical language for representing and manipulating uncertainty¹⁰, in much the same way as calculus is the language for representing and manipulating rates of change. Fortunately, the probabilistic approach to modelling is conceptually very simple: probability distributions are used to represent all the uncertain unobserved quantities in a model (including structural, parametric and noise-related) and how they relate to the data. Then the basic rules of probability theory are used to infer the unobserved quantities given the observed data. Learning from data occurs through the transformation of the prior probability distributions (defined before observing the data), into posterior distributions (after observing data). The application of probability theory to learning from data is called Bayesian learning (Box 1).

Apart from its conceptual simplicity, there are several appealing properties of the probabilistic framework for machine intelligence. Simple probability distributions over single or a few variables can be composed to form the building blocks of larger, more complex models. The dominant paradigm in machine learning over the past two decades for representing such compositional probabilistic models has been graphical models¹¹, with variants including directed graphs (also known as Bayesian networks and belief networks), undirected graphs (also known as Markov networks and random fields), and mixed graphs with both directed and undirected edges (Fig. 1). As discussed later, probabilistic programming offers an elegant way of generalizing graphical models, allowing a much richer representation of models. The compositionality of probabilistic models means that the behaviour of these building blocks in the context of the larger model is often much easier to understand than, say, what will happen if one couples a non-linear dynamical system (for example, a recurrent neural network) to another. In particular, for a well-defined probabilistic model, it is always possible to generate data from the model; such ‘imaginary’ data provide a window into the ‘mind’ of the probabilistic model, helping us to understand both the initial prior assumptions and what the model has learned at any later stage.

Probabilistic modelling also has some conceptual advantages over alternatives because it is a normative theory for learning in artificially intelligent systems. How should an artificially intelligent system represent and update its beliefs about the world in light of data? The Cox axioms define some desiderata for representing beliefs; a consequence of these axioms is that ‘degrees of belief’, ranging from ‘impossible’ to ‘absolutely certain’, must follow all the rules of probability theory^{10,12,13}. This justifies the use of subjective Bayesian probabilistic representations in artificial intelligence. An argument for Bayesian representations in artificial intelligence that is motivated by decision theory is given by the Dutch book theorem. The argument rests on the idea that the strength of beliefs of an agent can be assessed by asking the agent whether it would be willing to accept bets at various odds (ratios of payoffs). The Dutch book theorem states that unless an artificial intelligence system’s (or human’s, for that matter) degrees of beliefs are consistent with the rules of probability it will be willing to accept bets that are guaranteed to lose money¹⁴. Because of the force of these and many other arguments on the importance of a principled handling of uncertainty for intelligence, Bayesian probabilistic modelling has emerged not only as the theoretical foundation for rationality in artificial intelligence systems, but also as a model for normative behaviour in humans and animals^{15–18} (but see refs 19, 20 for a discussion), and much research is devoted to understanding how neural circuitry may be implementing Bayesian inference^{21,22}.

Although conceptually simple, a fully probabilistic approach to machine learning poses a number of computational and modelling challenges. Computationally, the main challenge is that learning involves marginalizing (summing out) all the variables in the model except for the variables of interest (Box 1). Such high-dimensional sums and integrals are generally computationally hard, in the sense that for many models

BOX 1

Bayesian machine learning

There are two simple rules that underlie probability theory: the sum rule:

$$P(x) = \sum_{y \in Y} P(x,y)$$

and the product rule:

$$P(x,y) = P(x)P(y|x).$$

Here x and y correspond to observed or uncertain quantities, taking values in some sets X and Y , respectively. For example, x and y might relate to the weather in Cambridge and London, respectively, both taking values in the set $X=Y=\{\text{rainy, cloudy, sunny}\}$. $P(x)$ corresponds to the probability of x , which can be either a statement about the frequency of observing a particular value, or a subjective belief about it. $P(x,y)$ is the joint probability of observing x and y , and $P(y|x)$ is the probability of y conditioned on observing the value of x . The sum rule states that the marginal of x is obtained by summing (or integrating for continuous variables) the joint over y . The product rule states that the joint can be decomposed as the product of the marginal and the conditional. Bayes rule is a corollary of these two rules:

$$P(y|x) = \frac{P(x|y)P(y)}{P(x)} = \frac{P(x|y)P(y)}{\sum_{y \in Y} P(x,y)}$$

We can apply probability theory to machine learning by replacing the symbols above: we replace x by D to denote the observed data, we replace y by θ to denote the unknown parameters of a model, and we condition all terms on m , the class of probabilistic models we are considering. For learning, we thus get:

$$P(\theta|D,m) = \frac{P(D|\theta,m)P(\theta|m)}{P(D|m)}$$

where $P(D|\theta,m)$ is the likelihood of parameters θ in model m , $P(\theta|m)$ is the prior probability of θ and $P(\theta|D,m)$ is the posterior of θ given data D .

For example, the data D might be a time series of hourly observations of the weather in Cambridge and London, and the model might attempt to capture the joint weather patterns at both locations over successive hours, with parameters θ modelling correlations over time and space. Learning is the transformation of prior knowledge or assumptions about the parameters $P(\theta|m)$, through the data D , into posterior knowledge about the parameters, $P(\theta|D,m)$. This posterior is now the prior to be used for future data. A learned model can be used to predict or forecast new unseen test data, D_{test} , by simply applying the sum and product rule to get the prediction:

$$P(D_{\text{test}}|D,m) = \int P(D_{\text{test}}|\theta,D,m)P(\theta|D,m)d\theta$$

Finally, different models can be compared by applying Bayes rule at the level of m :

$$P(m|D) = \frac{P(D|m)P(m)}{P(D)}$$

$$P(D|m) = \int P(D|\theta,m)P(\theta|m)d\theta$$

The term $P(D|m)$ is the marginal likelihood or model evidence, and implements a preference for simpler models known as Bayesian Ockham’s razor^{78,96,97}.

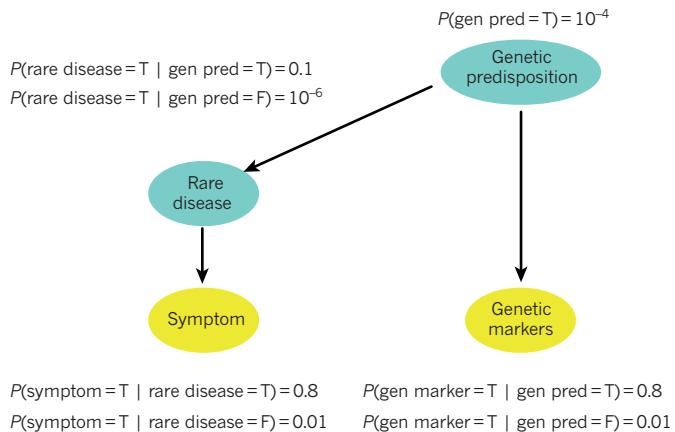


Figure 1 | Bayesian inference. A simple example of Bayesian inference applied to a medical diagnosis problem. Here the problem is diagnosing a rare disease using information from the patient's symptoms and, potentially, the patient's genetic marker measurements, which indicate predisposition (gen pred) to this disease. In this example, all variables are assumed to be binary. T, true; F, false. The relationships between variables are indicated by directed arrows and the probability of each variable given other variables they directly depend on is also shown. Yellow nodes denote measurable variables, whereas green nodes denote hidden variables. Using the sum rule (Box 1), the prior probability of the patient having the rare disease is: $P(\text{rare disease} = \text{T}) = P(\text{rare disease} = \text{T}|\text{gen pred} = \text{T})p(\text{gen pred} = \text{T}) + p(\text{rare disease} = \text{T}|\text{gen pred} = \text{F})p(\text{gen pred} = \text{F}) = 1.1 \times 10^{-5}$. Applying Bayes rule we find that for a patient observed to have the symptom, the probability of the rare disease is: $p(\text{rare disease} = \text{T}|\text{symptom} = \text{T}) = 8.8 \times 10^{-4}$, whereas for a patient observed to have the genetic marker (gen marker) it is $p(\text{rare disease} = \text{T}|\text{gen marker} = \text{T}) = 7.9 \times 10^{-4}$. Assuming that the patient has both the symptom and the genetic marker the probability of the rare disease increases to $p(\text{rare disease} = \text{T}|\text{symptom} = \text{T}, \text{gen marker} = \text{T}) = 0.06$. Here, we have shown fixed, known model parameters, that is, the numbers $\theta = (10^{-4}, 0.1, 10^{-6}, 0.8, 0.01, 0.8, 0.01)$. However, both these parameters and the structure of the model (the presence or absence of arrows and additional hidden variables) could be learned from a data set of patient records using the methods in Box 1.

there is no known polynomial time algorithm for performing them exactly. Fortunately, a number of approximate integration algorithms have been developed, including Markov chain Monte Carlo (MCMC) methods, variational approximations, expectation propagation and sequential Monte Carlo^{23–26}. It is worth noting that computational techniques are one area in which Bayesian machine learning differs from much of the rest of machine learning: for Bayesian researchers the main computational problem is integration, whereas for much of the rest of the community the focus is on optimization of model parameters. However, this dichotomy is not as stark as it appears: many gradient-based optimization methods can be turned into integration methods through the use of Langevin and Hamiltonian Monte Carlo methods^{27,28}, while integration problems can be turned into optimization problems through the use of variational approximations²⁴. I revisit optimization in a later section.

The main modelling challenge for probabilistic machine learning is that the model should be flexible enough to capture all the properties of the data required to achieve the prediction task of interest. One approach to addressing this challenge is to develop a prior distribution that encompasses an open-ended universe of models that can adapt in complexity to the data. The key statistical concept underlying flexible models that grow in complexity with the data is non-parametrics.

Flexibility through non-parametrics

One of the lessons of modern machine learning is that the best predictive performance is often obtained from highly flexible learning systems, especially when learning from large data sets. Flexible models can make better predictions because to a greater extent they allow data to ‘speak for themselves’. (But note that all predictions involve assumptions and

therefore the data are never exclusively ‘speaking for themselves.’) There are essentially two ways of achieving flexibility. The model could have a large number of parameters compared with the data set (for example, the neural network recently used to achieve translations of English and French sentences that approached the accuracy of state-of-the-art methods is a probabilistic model with 384 million parameters²⁹). Alternatively, the model can be defined using non-parametric components.

The best way to understand non-parametric models is through comparison to parametric ones. In a parametric model, there are a fixed, finite number of parameters, and no matter how much training data are observed, all the data can do is set these finitely many parameters that control future predictions. By contrast, non-parametric approaches have predictions that grow in complexity with the amount of training data, either by considering a nested sequence of parametric models with increasing numbers of parameters or by starting out with a model with infinitely many parameters. For example, in a classification problem, whereas a linear (parametric) classifier will always make predictions using a linear boundary between classes, a non-parametric classifier can learn a non-linear boundary whose shape becomes more complex with more data. Many non-parametric models can be derived starting from a parametric model and considering what happens as the model grows to the limit of infinitely many parameters³⁰. Clearly, fitting a model with infinitely many parameters to finite training data would result in ‘overfitting’, in the sense that the model’s predictions might reflect quirks of the training data rather than regularities that can be generalized to test data. Fortunately, Bayesian approaches are not prone to this kind of overfitting since they average over, rather than fit, the parameters (Box 1). Moreover, for many applications we have such huge data sets that the main concern is underfitting from the choice of an overly simplistic parametric model, rather than overfitting.

A full discussion of Bayesian non-parametrics is outside the scope of this Review (see refs 9, 31, 32 for this), but it is worth mentioning a few of the key models. Gaussian processes are a very flexible non-parametric model for unknown functions, and are widely used for regression, classification, and many other applications that require inference on functions³³. Consider learning a function that relates the dose of some chemical to the response of an organism to that chemical. Instead of modelling this relationship with, say, a linear parametric function, a Gaussian process could be used to learn directly a non-parametric distribution of non-linear functions consistent with the data. A notable example of a recent application of Gaussian processes is GaussianFace, a state-of-the-art approach to face recognition that outperforms humans and deep-learning methods³⁴. Dirichlet processes are a non-parametric model with a long history in statistics³⁵ and are used for density estimation, clustering, time-series analysis and modelling the topics of documents³⁶. To illustrate Dirichlet processes, consider an application to modelling friendships in a social network, where each person can belong to one of many communities. A Dirichlet process makes it possible to have a model whereby the number of inferred communities (that is, clusters) grows with the number of people³⁷. Dirichlet processes have also been used for clustering gene-expression patterns^{38,39}. The Indian buffet process (IBP)⁴⁰ is a non-parametric model that can be used for latent feature modelling, learning overlapping clusters, sparse matrix factorization, or to non-parametrically learn the structure of a deep network⁴¹. Elaborating the social network modelling example, an IBP-based model allows each person to belong to some subset of a large number of potential communities (for example, as defined by different families, workplaces, schools, hobbies, and so on) rather than a single community, and the probability of friendship between two people depends on the number of overlapping communities they have⁴². In this case, the latent features of each person correspond to the communities, which are not assumed to be observed directly. The IBP can be thought of as a way of endowing Bayesian non-parametric models with ‘distributed representations’, as popularized in the neural network literature⁴³. An interesting link between Bayesian non-parametrics and neural networks is that, under fairly general conditions, a neural network with

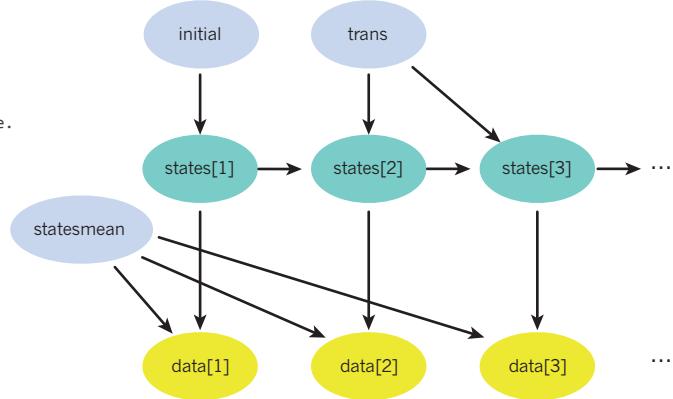
```

statesmean = [-1, 1, 0] # Emission parameters.
initial  = Categorical([1.0/3, 1.0/3, 1.0/3]) # Prob distr of state[1].
trans    = [Categorical([0.1, 0.5, 0.4]), Categorical([0.2, 0.2, 0.6]),
           Categorical([0.15, 0.15, 0.7])] # Trans distr for each state.
data     = [Nil, 0.9, 0.8, 0.7, 0, -0.025, -5, -2, -0.1, 0, 0.13]

@model hmm begin # Define a model hmm.
  states = Array(Int, length(data))
  @assume(states[1] ~ initial)
  for i = 2:length(data)
    @assume(states[i] ~ trans[states[i-1]])
    @observe(data[i] ~ Normal(statesmean[states[i]], 0.4))
  end
  @predict states
end

```

Figure 2 | Probabilistic programming. A probabilistic program in Julia (left) defining a simple three-state hidden Markov model (HMM), inspired by an example in ref. 62. The HMM is a widely used probabilistic model for sequential and time-series data, which assumes the data were obtained by transitioning stochastically between a discrete number of hidden states⁹⁸. The first four lines define the model parameters and the data. Here ‘trans’ is the 3×3 state-transition matrix, ‘initial’ is the initial state distribution, and ‘statesmean’ are the mean observations for each of the three states; actual observations are assumed to be noisy versions of this mean with Gaussian noise. The function hmm starts the definition of the HMM, drawing the



sequence of states with the @assume statements, and conditioning on the observed data with the @observe statements. Finally @predict states that we wish to infer the states and data; this inference is done automatically by the universal inference engine, which reasons over the configurations of this computer program. It would be trivial to modify this program so that the HMM parameters are unknown rather than fixed. A graphical model (right) corresponding to the HMM probabilistic program showing dependencies between the parameters (blue), hidden state variables (green) and observed data (yellow). This graphical model highlights the compositional nature of probabilistic models.

infinitely many hidden units is equivalent to a Gaussian process⁴⁴. Note that the above non-parametric components should be thought of again as building blocks, which can be composed into more complex models as described earlier. The next section describes an even more powerful way of composing models — through probabilistic programming.

Probabilistic programming

The basic idea in probabilistic programming is to use computer programs to represent probabilistic models (<http://probabilistic-programming.org>)^{45–47}. One way to do this is for the computer program to define a generator for data from the probabilistic model, that is, a simulator (Fig. 2). This simulator makes calls to a random number generator in such a way that repeated runs from the simulator would sample different possible data sets from the model. This simulation framework is more general than the graphical model framework described previously since computer programs can allow constructs such as recursion (functions calling themselves) and control flow statements (for example, ‘if’ statements that result in multiple paths a program can follow), which are difficult or impossible to represent in a finite graph. In fact, for many of the recent probabilistic programming languages that are based on extending Turing-complete languages (a class that includes almost all commonly used languages), it is possible to represent any computable probability distribution as a probabilistic program⁴⁸.

The full potential of probabilistic programming comes from automating the process of inferring unobserved variables in the model conditioned on the observed data (Box 1). Conceptually, conditioning needs to compute input states of the program that generate data matching the observed data. Whereas normally we think of programs running from inputs to outputs, conditioning involves solving the inverse problem of inferring the inputs (in particular the random number calls) that match a certain program output. Such conditioning is performed by a ‘universal inference engine’, usually implemented by Monte Carlo sampling over possible executions of the simulator program that are consistent with the observed data. The fact that defining such universal inference algorithms for computer programs is even possible is somewhat surprising, but it is related to the generality of certain key ideas from sampling such as rejection sampling, sequential Monte Carlo methods²³ and ‘approximate Bayesian computation’⁴⁹.

As an example, imagine you write a probabilistic program that simulates a gene regulatory model that relates unmeasured transcription

factors to the expression levels of certain genes. Your uncertainty in each part of the model would be represented by the probability distributions used in the simulator. The universal inference engine can then condition the output of this program on the measured expression levels, and automatically infer the activity of the unmeasured transcription factors and other uncertain model parameters. Another application of probabilistic programming implements a computer vision system as the inverse of a computer graphics program⁵⁰.

There are several reasons why probabilistic programming could prove to be revolutionary for machine intelligence and scientific modelling (its potential has been noticed by US Defense Advanced Research Projects Agency, which is currently funding a major programme called Probabilistic Programming for Advancing Machine Learning). First, the universal inference engine obviates the need to manually derive inference methods for models. Since deriving and implementing inference methods is generally the most rate-limiting and bug-prone step in modelling, often taking months, automating this step so that it takes minutes or seconds will greatly accelerate the deployment of machine learning systems. Second, probabilistic programming could be potentially transformative for the sciences, since it allows for rapid prototyping and testing of different models of data. Probabilistic programming languages create a very clear separation between the model and the inference procedures, encouraging model-based thinking⁵¹. There are a growing number of probabilistic programming languages. BUGS⁵², Stan⁵³, AutoBayes⁵⁴ and Infer.NET⁵⁵ allow only a restrictive class of models to be represented compared with systems based on Turing-complete languages. In return for this restriction, inference in such languages can be much faster than for the more general languages⁵⁶, such as IBAL⁵⁷, BLOG⁵⁸, Church⁵⁹, Figaro⁶⁰, Venture⁶¹, and Anglican⁶². A major emphasis of recent work is on fast inference in general languages (see for example ref. 63). Nearly all approaches to probabilistic programming are Bayesian since it is hard to create other coherent frameworks for automated reasoning about uncertainty. Notable exceptions are systems such as Theano, which is not itself a probabilistic programming language but uses symbolic differentiation to speed up and automate optimization of parameters of neural networks and other probabilistic models⁶⁴.

Although parameter optimization is commonly used to improve probabilistic models, in the next section I will describe recent work on how probabilistic modelling can be used to improve optimization.

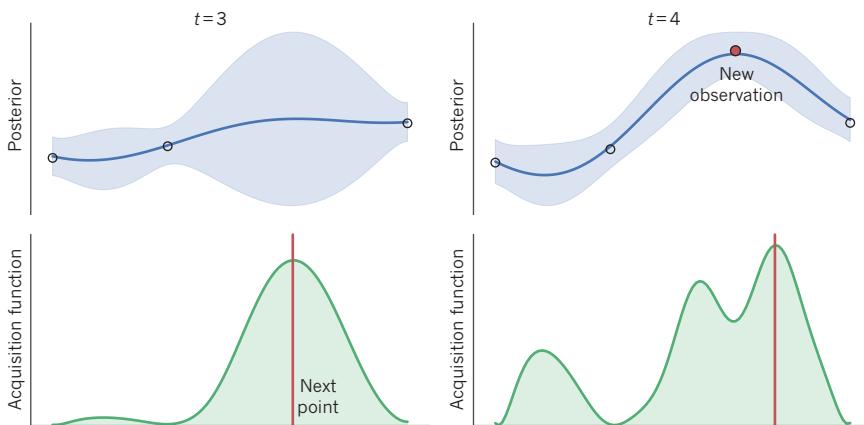


Figure 3 | Bayesian optimization. A simple illustration of Bayesian optimization in one dimension. The goal is to maximize some true unknown function f (not shown). Information about this function is gained by making observations (circles, top), which are evaluations of the function at specific x values. These observations are used to infer a posterior distribution over the function values (shown as mean, blue line; and standard deviations, blue shaded area) representing the distribution of possible functions; note that uncertainty grows away from the observations. On the basis of this distribution over functions, an acquisition function is

computed (green shaded area, bottom panels), which represents the gain from evaluating the unknown function f at different x values; note that the acquisition function is high where the posterior over f has both high mean and large uncertainty. Different acquisition functions can be used such as ‘expected improvement’ or ‘information gain’. The peak of the acquisition function (red line) is the best next point to evaluate, and is therefore chosen for evaluation (red dot, new observation). The left and right panels show an example of what could happen after three and four function evaluations, respectively.

Bayesian optimization

Consider the very general problem of finding the global maximum of an unknown function that is expensive to evaluate (say, evaluating the function requires performing lots of computation or conducting an experiment). Mathematically, for a function f on a domain X , the goal is to find a global maximizer x^* :

$$x^* = \arg \max_{x \in X} f(x).$$

Bayesian optimization poses this as a problem in sequential decision theory: where should one evaluate next so as most quickly to maximize f , taking into account the gain in information about the unknown function f (refs 65,66)? For example, having evaluated at three points, measuring the corresponding values of the function at those points, $[(x_1, f(x_1)), (x_2, f(x_2)), (x_3, f(x_3))]$, which point x should the algorithm evaluate next, and where does it believe the maximum to be? This is a classic machine-intelligence problem with a wide range of applications in science and engineering, from drug design (where the function could be the drug’s efficacy) to robotics (where the function could be the speed of a robot’s gait). It can be applied to any problem involving the optimization of expensive functions; the qualifier ‘expensive’ comes because Bayesian optimization might use substantial computational resources to decide where to evaluate next, and a trade-off for these resources has to be made with the cost of function evaluations.

The current best-performing global optimization methods maintain a Bayesian representation of the probability distribution over the uncertain function f being optimized, and use this uncertainty to decide where (in X) to query next^{67–69}. In continuous spaces, most Bayesian optimization methods (Fig. 3) use Gaussian processes (as described in the section on non-parametrics) to model the unknown function. A recent high-impact application has been the optimization of the training process for machine-learning models, including deep neural networks⁷⁰. This and related recent work⁷¹ are further examples of the application of machine intelligence to improve machine intelligence.

There are interesting links between Bayesian optimization and reinforcement learning. Specifically, Bayesian optimization is a sequential decision problem where the decisions (choices of x to evaluate) do not affect the state of the system (the actual function f). Such state-less sequential decision problems fall under the rubric of multi-arm bandits⁷².

a subclass of reinforcement-learning problems. More broadly, important recent work takes a Bayesian approach to learning to control uncertain systems⁷³ (for a review see ref. 74). Faithfully representing uncertainty about the future outcome of actions is particularly important in decision and control problems. Good decisions rely on good representations of the probability of different outcomes and their relative payoffs.

More generally, Bayesian optimization is a special case of Bayesian numerical computation^{75,76}, which is re-emerging as a very active area of research (<http://www.probabilistic-numerics.org>), and includes topics such as solving ordinary differential equations and numerical integration. In all these cases, probability theory is being used to represent computational uncertainty; that is, the uncertainty that one has about the outcome of a deterministic computation.

Data compression

Consider the problem of compressing data so as to communicate them or store them in as few bits as possible in such a manner that the original data can be recovered exactly from the compressed data. Methods for such lossless data compression are ubiquitous in information technology, from computer hard drives to data transfer over the internet. Data compression and probabilistic modelling are two sides of the same coin, and Bayesian machine-learning methods are increasingly advancing the state-of-the-art in compression. The connection between compression and probabilistic modelling was established in the mathematician Claude Shannon’s seminal work on the source coding theorem⁷⁷, which states that the number of bits required to compress data in a lossless manner is bounded by the entropy of the probability distribution of the data. All commonly used lossless data compression algorithms (for example, gzip) can be viewed as probabilistic models of sequences of symbols.

The link to Bayesian machine learning is that the better the probabilistic model one learns, the higher the compression rate can be⁷⁸. These models need to be flexible and adaptive, since different kinds of sequences have very different statistical patterns (say, Shakespeare’s plays or computer source code). It turns out that some of the world’s best compression algorithms (for example, Sequence Memoizer⁷⁹ and PPM with dynamic parameter updates⁸⁰) are equivalent to Bayesian non-parametric models of sequences, and improvements to compression are being made through a better understanding of how to learn the statistical structure of sequences. Future advances in compression will come with advances in

probabilistic machine learning, including special compression methods for non-sequence data such as images, graphs and other structured objects.

Automatic discovery of interpretable models from data

One of the grand challenges of machine learning is to fully automate the process of learning and explaining statistical models from data. This is the goal of the Automatic Statistician (<http://www.automaticstatistician.com>), a system that can automatically discover plausible models from data, and explain what it has discovered in plain English⁸¹. This could be useful to almost any field of endeavour that is reliant on extracting knowledge from data. In contrast to the methods described in much of the machine-learning literature, which have been focused on extracting increasing

performance improvements on pattern-recognition problems using techniques such as kernel methods, random forests or deep learning, the Automatic Statistician builds models that are composed of interpretable components, and has a principled way of representing uncertainty about model structures given the data. It also gives reasonable answers not just for big data sets, but also for small ones. Bayesian approaches provide an elegant way of trading off the complexity of the model and the complexity of the data, and probabilistic models are compositional and interpretable, as already described.

A prototype version of the Automatic Statistician takes in time-series data and automatically generates 5–15 page reports describing the model it has discovered (Fig. 4). This system is based on the idea that probabilistic building blocks can be combined through a grammar to build

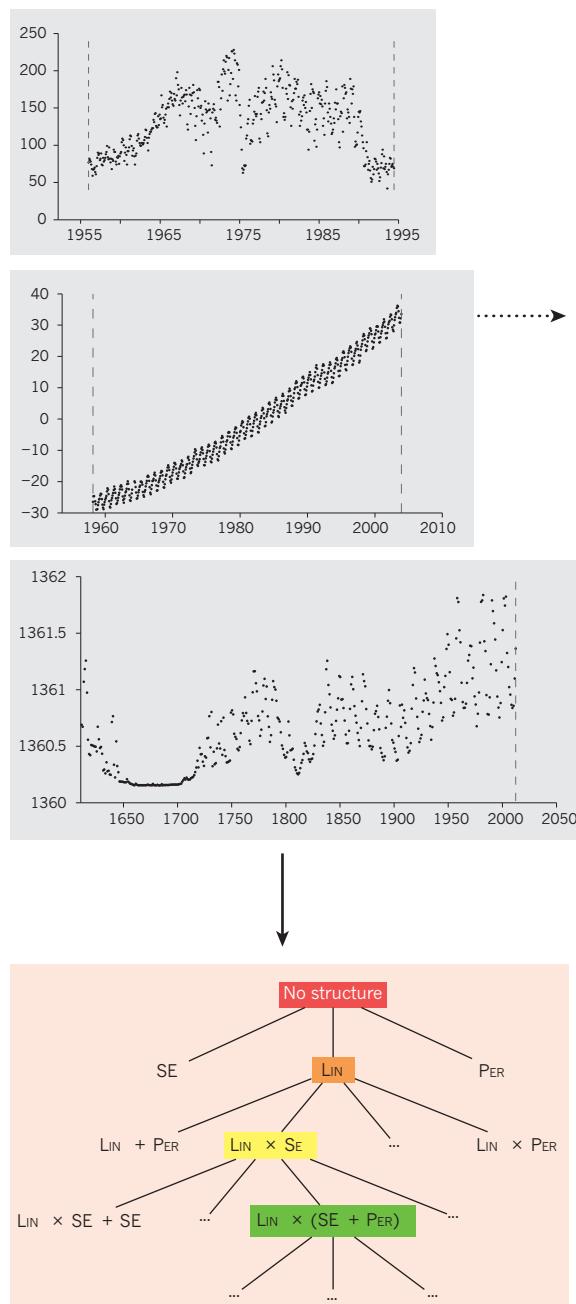


Figure 4 | The Automatic Statistician. A flow diagram describing the Automatic Statistician. The input to the system is data, in this case represented as time series (top left). The system searches over a grammar of models to discover a good interpretation of the data (bottom left), using Bayesian inference to score the models (Box 1). Components of the model

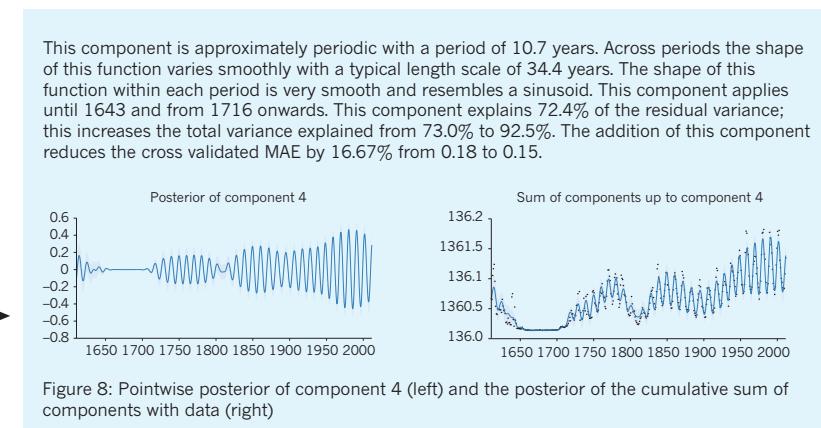
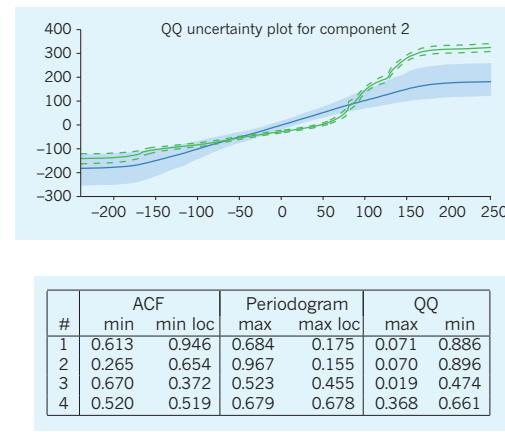


Figure 8: Pointwise posterior of component 4 (left) and the posterior of the cumulative sum of components with data (right)



discovered are translated into English phrases (bottom right). The end result is a report with text, figures and tables, describing in detail what has been inferred about the data, including a section on model checking and criticism (top right)^{99,100}. Data and the report are for illustrative purposes only.

an open-ended language of models⁸². In contrast to work on equation learning (see for example ref. 83), the models attempt to capture general properties of functions (for example, smoothness, periodicity or trends) rather than a precise equation. Handling uncertainty is at the core of the Automatic Statistician; it makes use of Bayesian non-parametrics to give it the flexibility to obtain state-of-the-art predictive performance, and uses the metric marginal likelihood (Box 1) to search the space of models.

Important earlier work includes statistical expert systems^{84,85} and the Robot Scientist, which integrates machine learning and scientific discovery in a closed loop with an experimental platform in microbiology to automate the design and execution of new experiments⁸⁶. Auto-WEKA is a recent project that automates learning classifiers, making heavy use of the Bayesian optimization techniques already described⁷¹. Efforts to automate the application of machine-learning methods to data have recently gained momentum, and may ultimately result in artificial intelligence systems for data science.

Perspective

The information revolution has resulted in the availability of ever larger collections of data. What is the role of uncertainty in modelling such big data? Classic statistical results state that under certain regularity conditions, in the limit of large data sets the posterior distribution of the parameters for Bayesian parametric models converges to a single point around the maximum likelihood estimate. Does this mean that Bayesian probabilistic modelling of uncertainty is unnecessary if you have a lot of data?

There are at least two reasons this is not the case⁸⁷. First, as we have seen, Bayesian non-parametric models have essentially infinitely many parameters, so no matter how many data one has, their capacity to learn should not saturate, and their predictions should continue to improve.

Second, many large data sets are in fact large collections of small data sets. For example, in areas such as personalized medicine and recommendation systems, there might be a large amount of data, but there is still a relatively small amount of data for each patient or client, respectively. To customize predictions for each person it becomes necessary to build a model for each person — with its inherent uncertainties — and to couple these models together in a hierarchy so that information can be borrowed from other similar people. We call this the personalization of models, and it is naturally implemented using hierarchical Bayesian approaches such as hierarchical Dirichlet processes³⁶, and Bayesian multitask learning^{88,89}.

Probabilistic approaches to machine learning and intelligence are a very active area of research with wide-ranging impact beyond conventional pattern-recognition problems. As I have outlined, these problems include data compression, optimization, decision making, scientific model discovery and interpretation, and personalization. The key distinction between problems in which a probabilistic approach is important and problems that can be solved using non-probabilistic machine-learning approaches is whether uncertainty has a central role. Moreover, most conventional optimization-based machine-learning approaches have probabilistic analogues that handle uncertainty in a more principled manner. For example, Bayesian neural networks represent the parameter uncertainty in neural networks⁴⁴, and mixture models are a probabilistic analogue for clustering methods⁷⁸. Although probabilistic machine learning often defines how to solve a problem in principle, the central challenge in the field is finding how to do so in practice in a computationally efficient manner^{90,91}. There are many approaches to the efficient approximation of computationally hard inference problems. Modern inference methods have made it possible to scale to millions of data points, making probabilistic methods computationally competitive with conventional methods^{92–95}. Ultimately, intelligence relies on understanding and acting in an imperfectly sensed and uncertain world. Probabilistic modelling will continue to play a central part in the development of ever more powerful machine learning and artificial intelligence systems. ■

Received 12 February; accepted 21 April 2015.

- Russell, S. & Norvig, P. *Artificial Intelligence: a Modern Approach* (Prentice-Hall, 1995).
- Thrun, S., Burgard, W. & Fox, D. *Probabilistic Robotics* (MIT Press, 2006).

- Bishop, C. M. *Pattern Recognition and Machine Learning* (Springer, 2006).
- Murphy, K. P. *Machine Learning: A Probabilistic Perspective* (MIT Press, 2012).
- Hinton, G. et al. Deep neural networks for acoustic modeling in speech recognition: the shared views of four research groups. *IEEE Signal Process. Mag.* **29**, 82–97 (2012).
- Krizhevsky, A., Sutskever, I. & Hinton, G. E. ImageNet classification with deep convolutional neural networks. In *Proc. Advances in Neural Information Processing Systems 25* 1097–1105 (2012).
- Sermanet, P. et al. Overfeat: integrated recognition, localization and detection using convolutional networks. In *Proc. International Conference on Learning Representations* <http://arxiv.org/abs/1312.6229> (2014).
- Bengio, Y., Ducharme, R., Vincent, P. & Janvin, C. A neural probabilistic language model. *J. Mach. Learn. Res.* **3**, 1137–1155 (2003).
- Ghahramani, Z. Bayesian nonparametrics and the probabilistic approach to modelling. *Phil. Trans. R. Soc. A* **371**, 20110553 (2013). **A review of Bayesian non-parametric modelling written for a general scientific audience.**
- Jaynes, E. T. *Probability Theory: the Logic of Science* (Cambridge Univ. Press, 2003).
- Koller, D. & Friedman, N. *Probabilistic Graphical Models: Principles and Techniques* (MIT Press, 2009). **This is an encyclopaedic text on probabilistic graphical models spanning many key topics.**
- Cox, R. T. *The Algebra of Probable Inference* (Johns Hopkins Univ. Press, 1961).
- Van Horn, K. S. Constructing a logic of plausible inference: a guide to Cox's theorem. *Int. J. Approx. Reason.* **34**, 3–24 (2003).
- De Finetti, B. La prévision: ses lois logiques, ses sources subjectives. In *Annales de l'institut Henri Poincaré* [in French] **7**, 1–68 (1937).
- Knill, D. & Richards, W. Perception as Bayesian inference (Cambridge Univ. Press, 1996).
- Griffiths, T. L. & Tenenbaum, J. B. Optimal predictions in everyday cognition. *Psychol. Sci.* **17**, 767–773 (2006).
- Wolpert, D. M., Ghahramani, Z. & Jordan, M. I. An internal model for sensorimotor integration. *Science* **269**, 1880–1882 (1995).
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L. & Goodman, N. D. How to grow a mind: statistics, structure, and abstraction. *Science* **331**, 1279–1285 (2011).
- Marcus, G. F. & Davis, E. How robust are probabilistic models of higher-level cognition? *Psychol. Sci.* **24**, 2351–2360 (2013).
- Goodman, N. D. et al. Relevant and robust a response to Marcus and Davis (2013). *Psychol. Sci.* **26**, 539–541 (2015).
- Doya, K., Ishii, S., Pouget, A. & Rao, R. P. N. *Bayesian Brain: Probabilistic Approaches to Neural Coding* (MIT Press, 2007).
- Deneve, S. Bayesian spiking neurons I: inference. *Neural Comput.* **20**, 91–117 (2008).
- Neal, R. M. *Probabilistic Inference Using Markov Chain Monte Carlo Methods*. Report No. CRG-TR-93-1 <http://www.cs.toronto.edu/~radford/review.abstract.html> (Univ. Toronto, 1993).
- Jordan, M., Ghahramani, Z., Jaakkola, T. & Saul, L. An introduction to variational methods in graphical models. *Mach. Learn.* **37**, 183–233 (1999).
- Doucet, A., de Freitas, J. F. G. & Gordon, N. J. *Sequential Monte Carlo Methods in Practice* (Springer, 2000).
- Minka, T. P. Expectation propagation for approximate Bayesian inference. In *Proc. Uncertainty in Artificial Intelligence 17* 362–369 (2001).
- Neal, R. M. In *Handbook of Markov Chain Monte Carlo* (eds Brooks, S., Gelman, A., Jones, G. & Meng, X.-L.) (Chapman & Hall/CRC, 2010).
- Girolami, M. & Calderhead, B. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *J. R. Stat. Soc. Series B Stat. Methodol.* **73**, 123–214 (2011).
- Sutskever, I., Vinyals, O. & Le, Q. V. Sequence to sequence learning with neural networks. In *Proc. Advances in Neural Information Processing Systems 27*, 3104–3112 (2014).
- Neal, R. M. In *Maximum Entropy and Bayesian Methods* 197–211 (Springer, 1992).
- Orbanz, P. & Teh, Y. W. in *Encyclopedia of Machine Learning* 81–89 (Springer, 2010).
- Hjort, N., Holmes, C., Müller, P. & Walker, S. (eds). *Bayesian Nonparametrics* (Cambridge Univ. Press, 2010).
- Rasmussen, C. E. & Williams, C. K. I. *Gaussian Processes for Machine Learning* (MIT Press, 2006). **This is a classic monograph on Gaussian processes, relating them to kernel methods and other areas of machine learning.**
- Lu, C. & Tang, X. Surpassing human-level face verification performance on LFW with GaussianFace. In *Proc. 29th AAAI Conference on Artificial Intelligence* <http://arxiv.org/abs/1404.3840> (2015).
- Ferguson, T. S. A Bayesian analysis of some nonparametric problems. *Ann. Stat.* **1**, 209–230 (1973).
- Teh, Y. W., Jordan, M. I., Beal, M. J. & Blei, D. M. Hierarchical Dirichlet processes. *J. Am. Stat. Assoc.* **101**, 1566–1581 (2006).
- Kemp, C., Tenenbaum, J. B., Griffiths, T. L., Yamada, T. & Ueda, N. Learning systems of concepts with an infinite relational model. In *Proc. 21st National Conference on Artificial Intelligence* 381–388 (2006).
- Medvedovic, M. & Sivaganesan, S. Bayesian infinite mixture model based clustering of gene expression profiles. *Bioinformatics* **18**, 1194–1206 (2002).
- Rasmussen, C. E., De la Cruz, B. J., Ghahramani, Z. & Wild, D. L. Modeling and visualizing uncertainty in gene expression clusters using Dirichlet process mixtures. *Trans. Comput. Biol. Bioinform.* **6**, 615–628 (2009).

40. Griffiths, T. L. & Ghahramani, Z. The Indian buffet process: an introduction and review. *J. Mach. Learn. Res.* **12**, 1185–1224 (2011).
This article introduced a new class of Bayesian non-parametric models for latent feature modelling.

41. Adams, R. P., Wallach, H. & Ghahramani, Z. Learning the structure of deep sparse graphical models. In *Proc. 13th International Conference on Artificial Intelligence and Statistics* (eds Teh, Y. W. & Titterington, M.) 1–8 (2010).

42. Miller, K., Jordan, M. I. & Griffiths, T. L. Nonparametric latent feature models for link prediction. In *Proc. Advances in Neural Information Processing Systems* 1276–1284 (2009).

43. Hinton, G. E., McClelland, J. L. & Rumelhart, D. E. in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Foundations* 77–109 (MIT Press, 1986).

44. Neal, R. M. *Bayesian Learning for Neural Networks* (Springer, 1996).
This text derived MCMC-based Bayesian inference in neural networks and drew important links to Gaussian processes.

45. Koller, D., McAllester, D. & Pfeffer, A. Effective Bayesian inference for stochastic programs. In *Proc. 14th National Conference on Artificial Intelligence* 740–747 (1997).

46. Goodman, N. D. & Stuhlmüller, A. *The Design and Implementation of Probabilistic Programming Languages*. Available at <http://dippl.org> (2015).

47. Pfeffer, A. *Practical Probabilistic Programming* (Manning, 2015).

48. Freer, C., Roy, D. & Tenenbaum, J. B. in *Turing's Legacy* (ed. Downey, R.) 195–252 (2014).

49. Marjoram, P., Molitor, J., Plagnol, V. & Tavaré, S. Markov chain Monte Carlo without likelihoods. *Proc. Natl Acad. Sci. USA* **100**, 15324–15328 (2003).

50. Mansinghka, V., Kulkarni, T. D., Perov, Y. N. & Tenenbaum, J. Approximate Bayesian image interpretation using generative probabilistic graphics programs. In *Proc. Advances in Neural Information Processing Systems* 26 1520–1528 (2013).

51. Bishop, C. M. Model-based machine learning. *Phil. Trans. R. Soc. A* **371**, 20120222 (2013).
This article is a very clear tutorial exposition of probabilistic modelling.

52. Lunn, D. J., Thomas, A., Best, N. & Spiegelhalter, D. WinBUGS — a Bayesian modelling framework: concepts, structure, and extensibility. *Stat. Comput.* **10**, 325–337 (2000).
This reports an early probabilistic programming framework widely used in statistics.

53. Stan Development Team. *Stan Modeling Language Users Guide and Reference Manual, Version 2.5.0*. <http://mc-stan.org/> (2014).

54. Fischer, B. & Schumann, J. AutoBayes: a system for generating data analysis programs from statistical models. *J. Funct. Program.* **13**, 483–508 (2003).

55. Minka, T. P., Winn, J. M., Guiver, J. P. & Knowles, D. A. *Infer.NET* 2.4. <http://research.microsoft.com/infernet> (Microsoft Research, 2010).

56. Wingate, D., Stuhlmüller, A. & Goodman, N. D. Lightweight implementations of probabilistic programming languages via transformational compilation. In *Proc. International Conference on Artificial Intelligence and Statistics* 770–778 (2011).

57. Pfeffer, A. IBAL: a probabilistic rational programming language. In *Proc. International Joint Conference on Artificial Intelligence* 733–740 (2001).

58. Milch, B. et al. BLOG: probabilistic models with unknown objects. In *Proc. 19th International Joint Conference on Artificial Intelligence* 1352–1359 (2005).

59. Goodman, N., Mansinghka, V., Roy, D., Bonawitz, K. & Tenenbaum, J. Church: a language for generative models. In *Proc. Uncertainty in Artificial Intelligence* 22 23 (2008).
This is an influential paper introducing the Turing-complete probabilistic programming language Church.

60. Pfeffer, A. *Figaro: An Object-Oriented Probabilistic Programming Language*. Tech. Rep. (Charles River Analytics, 2009).

61. Mansinghka, V., Selsam, D. & Perov, Y. Venture: a higher-order probabilistic programming platform with programmable inference. Preprint at <http://arxiv.org/abs/1404.0099> (2014).

62. Wood, F., van de Meent, J. W. & Mansinghka, V. A new approach to probabilistic programming inference. In *Proc. 17th International Conference on Artificial Intelligence and Statistics* 1024–1032 (2014).

63. Li, L., Wu, Y. & Russell, J. S. *SWIFT: Compiled Inference for Probabilistic Programs*. Report No. UCB/EECS-2015-12 (Univ. California, Berkeley, 2015).

64. Bergstra, J. et al. Theano: a CPU and GPU math expression compiler. In *Proc. 9th Python in Science Conference* <http://conference.scipy.org/proceedings/scipy2010/> (2010).

65. Kushner, H. A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. *J. Basic Eng.* **86**, 97–106 (1964).

66. Jones, D. R., Schonlau, M. & Welch, W. J. Efficient global optimization of expensive black-box functions. *J. Glob. Optim.* **13**, 455–492 (1998).

67. Brochu, E., Cora, V. M. & de Freitas, N. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. Preprint at <http://arXiv.org/abs/1012.2599> (2010).

68. Hennig, P. & Schuler, C. J. Entropy search for information-efficient global optimization. *J. Mach. Learn. Res.* **13**, 1809–1837 (2012).

69. Hernández-Lobato, J. M., Hoffman, M. W. & Ghahramani, Z. Predictive entropy search for efficient global optimization of black-box functions. In *Proc. Advances in Neural Information Processing Systems* 918–926 (2014).

70. Snoek, J., Larochelle, H. & Adams, R. P. Practical Bayesian optimization of machine learning algorithms. In *Proc. Advances in Neural Information Processing Systems* 2960–2968 (2012).

71. Thornton, C., Hutter, F., Hoos, H. H. & Leyton-Brown, K. Auto-WEKA: combined selection and hyperparameter optimization of classification algorithms. In *Proc. 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 847–855 (2013).

72. Robbins, H. Some aspects of the sequential design of experiments. *Bull. Amer. Math. Soc.* **55**, 527–535 (1952).

73. Deisenroth, M. P. & Rasmussen, C. E. PILCO: a model-based and data-efficient approach to policy search. In *Proc. 28th International Conference on Machine Learning* 465–472 (2011).

74. Poupart, P. in *Encyclopedia of Machine Learning* 90–93 (Springer, 2010).

75. Diaconis, P. in *Statistical Decision Theory and Related Topics IV* 163–175 (Springer, 1988).

76. O'Hagan, A. Bayes-Hermite quadrature. *J. Statist. Plann. Inference* **29**, 245–260 (1991).

77. Shannon, C. & Weaver, W. *The Mathematical Theory of Communication* (Univ. Illinois Press, 1949).

78. MacKay, D. J. C. *Information Theory, Inference, and Learning Algorithms* (Cambridge Univ. Press, 2003).

79. Wood, F., Gasthaus, J., Archambeau, C., James, L. & Teh, Y. W. The sequence memoizer. *Commun. ACM* **54**, 91–98 (2011).
This article derives a state-of-the-art data compression scheme based on Bayesian nonparametric models.

80. Steinruecken, C., Ghahramani, Z. & MacKay, D. J. C. Improving PPM with dynamic parameter updates. In *Proc. Data Compression Conference* (in the press).

81. Lloyd, J. R., Duvenaud, D., Grosse, R., Tenenbaum, J. B. & Ghahramani, Z. Automatic construction and natural-language description of nonparametric regression models. In *Proc. 28th AAAI Conference on Artificial Intelligence* Preprint at: <http://arxiv.org/abs/1402.4304> (2014).
Introduces the Automatic Statistician, translating learned probabilistic models into reports about data.

82. Grosse, R. B., Salakhutdinov, R. & Tenenbaum, J. B. Exploiting compositionality to explore a large space of model structures. In *Proc. Conference on Uncertainty in Artificial Intelligence* 306–315 (2012).

83. Schmidt, M. & Lipson, H. Distilling free-form natural laws from experimental data. *Science* **324**, 81–85 (2009).

84. Wolstenholme, D. E., O'Brien, C. M. & Nelder, J. A. GLIMPSE: a knowledge-based front end for statistical analysis. *Knowl. Base. Syst.* **1**, 173–178 (1988).

85. Hand, D. J. Patterns in statistical strategy. In *Artificial Intelligence and Statistics* (ed Gale, W. A.) (Addison-Wesley Longman, 1986).

86. King, R. D. et al. Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature* **427**, 247–252 (2004).

87. Welling, M. et al. Bayesian inference with big data: a snapshot from a workshop. *ISBA Bulletin* **21**, <https://bayesian.org/sites/default/files/fm/bulletins/1412.pdf> (2014).

88. Bakker, B. & Heskes, T. Task clustering and gating for Bayesian multitask learning. *J. Mach. Learn. Res.* **4**, 83–99 (2003).

89. Houlsby, N., Hernández-Lobato, J. M., Huszár, F. & Ghahramani, Z. Collaborative Gaussian processes for preference learning. In *Proc. Advances in Neural Information Processing Systems* 26 2096–2104 (2012).

90. Russell, S. J. & Wefald, E. *Do the Right Thing: Studies in Limited Rationality* (MIT Press, 1991).

91. Jordan, M. I. On statistics, computation and scalability. *Bernoulli* **19**, 1378–1390 (2013).

92. Hoffman, M., Blei, D., Paisley, J. & Wang, C. Stochastic variational inference. *J. Mach. Learn. Res.* **14**, 1303–1347 (2013).

93. Hensman, J., Fusi, N. & Lawrence, N. D. Gaussian processes for big data. In *Proc. Conference on Uncertainty in Artificial Intelligence* 244 (UAI, 2013).

94. Korattikara, A., Chen, Y. & Welling, M. Austerity in MCMC land: cutting the Metropolis-Hastings budget. In *Proc. 31st International Conference on Machine Learning* 181–189 (2014).

95. Paige, B., Wood, F., Doucet, A. & Teh, Y. W. Asynchronous anytime sequential Monte Carlo. In *Proc. Advances in Neural Information Processing Systems* 27 3410–3418 (2014).

96. Jefferys, W. H. & Berger, J. O. Ockham's Razor and Bayesian Analysis. *Am. Sci.* **80**, 64–72 (1992).

97. Rasmussen, C. E. & Ghahramani, Z. Occam's Razor. In *Neural Information Processing Systems* 13 (eds Leen, T. K., Dietterich, T. G., & Tresp, V.) 294–300 (2001).

98. Rabiner, L. R. A tutorial on hidden Markov models and selected applications in speech recognition. *Proc. IEEE* **77**, 257–286 (1989).

99. Gelman, A. et al. *Bayesian Data Analysis* 3rd edn (Chapman & Hall/CRC, 2013).

100. Lloyd, J. R. & Ghahramani, Z. Statistical model criticism using kernel two sample tests (2015).

Acknowledgements I acknowledge an EPSRC grant EP/I036575/1, the DARPA PPAML programme, a Google Focused Research Award for the Automatic Statistician and support from Microsoft Research. I am also grateful for valuable input from D. Duvenaud, H. Ge, M. W. Hoffman, J. R. Lloyd, A. Ścibior, M. Welling and D. Wolpert.

Author Information Reprints and permissions information is available at [www.nature.com/reprints](http://nature.com/reprints). The author declares competing financial interests: see go.nature.com/hyx7hh for details. Readers are welcome to comment on the online version of this paper at go.nature.com/hyx7hh. Correspondence should be addressed to Z.G. (zoubin@eng.cam.ac.uk).

Science, technology and the future of small autonomous drones

Dario Floreano¹ & Robert J. Wood²

We are witnessing the advent of a new era of robots — drones — that can autonomously fly in natural and man-made environments. These robots, often associated with defence applications, could have a major impact on civilian tasks, including transportation, communication, agriculture, disaster mitigation and environment preservation. Autonomous flight in confined spaces presents great scientific and technical challenges owing to the energetic cost of staying airborne and to the perceptual intelligence required to negotiate complex environments. We identify scientific and technological advances that are expected to translate, within appropriate regulatory frameworks, into pervasive use of autonomous drones for civilian applications.

Drones, a popular nickname for unmanned aerial vehicles and micro aerial vehicles, often conjure up images of unmanned aeroplanes that fly thousands of miles for espionage and to deploy munitions. However, over the past few years, an increasing number of public and private research laboratories have been working on small, human-friendly drones that one day may autonomously fly in confined spaces and in close proximity to people. The development of these small drones, which is the main focus of this Review, has been supported by the miniaturization and cost reduction of electronic components (micro-processors, sensors, batteries and wireless communication units), largely driven by the portable electronic device industry. These improvements have enabled the prototyping and commercialization of small (typically less than 1 kg) drones at smartphone prices.

Small drones will have important socio-economic impacts (Fig. 1). Images from drones that are capable of flying a few metres above the ground will fill a gap between expensive, weather-dependent and low-resolution images provided by satellites and car-based images limited to human-level perspectives and the availability of accessible roads. Specialized flying cameras and cloud-based data analytics will allow farmers to continuously monitor the quality of crop growth. Such platforms will enable construction companies to measure work progress in real time. Drones will let mining companies obtain precise volumetric data of excavations. Energy and infrastructure companies will be able to exhaustively survey pipelines, roads and cables. Humanitarian organizations could immediately assess and adapt aid efforts in continuously changing refugee camps. Transportation drones that are capable of safely taking off and landing in the proximity of buildings and humans will allow developing countries — without a suitable road network — to rapidly deliver goods and to finally unleash the full potential of their e-commerce telecommunication infrastructure. Transportation drones will also help developed countries to improve the quality of service in congested or remote areas, and will enable rescue organizations to quickly deliver medical supplies in the field and on demand. Inspection drones that are capable of flying in confined spaces will help fire-fighting and emergency units to assess dangers faster and more safely, logistic companies to detect cracks in the inner and outer shells of ships, road maintenance companies to measure signs of wear and tear in bridges and tunnels, security companies to improve building safety by monitoring areas outside the range of surveillance cameras, and disaster mitigation agencies to inspect partially collapsed

buildings where ground clutter is an obstacle for terrestrial robots. Coordinated teams of autonomous drones will enable missions that last longer than the flight time of a single drone by allowing some drones to temporarily leave the team for battery replacement. Drone teams will permit rescue organizations to quickly deploy dedicated communication networks for ground operators. Telecommunication companies can also leverage drone networks to temporarily supplement or replace points of service.

Realizing this vision will require new mechatronic solutions and new levels of control autonomy (Box 1) to safely complete missions in confined spaces and near the ground, where absolute positioning signals and remote control are not available or sufficiently precise. None of today's commercial drones has sufficient control autonomy to complete any of the missions described without skilled human supervision, which makes those operations slow, dangerous and not scalable. We identify the most promising scientific and technological advances that could lead to a new generation of small autonomous drones and offer a tentative road map of capability deployment within suitable regulatory frameworks.

Design and manufacturing challenges

Flight is energetically expensive, particularly when the size of the device is reduced. This is often due to practical issues that arise when scaling a vehicle down, such as reduced power density of electromagnetic motors, decreased transmission efficiency owing to increased dominance of friction from gears and bearings, and greater viscous losses because of reduced Reynolds numbers. Certain flight modes are also challenging — for example, hovering — either owing to energetics or control; these also become increasingly difficult as size is reduced. Scaling issues create problems at the most basic level of autonomy: they limit the ability to simply sustain flight for an adequate amount of time to perform higher-level mission functions. To overcome these challenges, developers of drones must consider system-level design choices that balance trade-offs that arise when selecting the constituent components to power the vehicle. The most common morphologies for flying robots include more conventional fixed-wing and rotary-wing aircraft (rotorcraft) designs as well as bio-inspired designs based on flapping wings. Each of these platform types has pros and cons. For example, fixed-wing aircraft (Fig. 2a) are capable of fast and efficient flight, but typically cannot hover. Rotorcraft (Fig. 2b) can hover and are highly manoeuvrable, but are generally less efficient in forward flight than fixed-wing vehicles. Neither fixed-wing

¹Laboratory of Intelligent Systems, Ecole Polytechnique Fédérale de Lausanne, Station 11, Lausanne, CH 1015, Switzerland. ²Microrobotics Lab, Harvard University, 149 Maxwell Dworkin Building, Cambridge, Massachusetts 02138, USA.

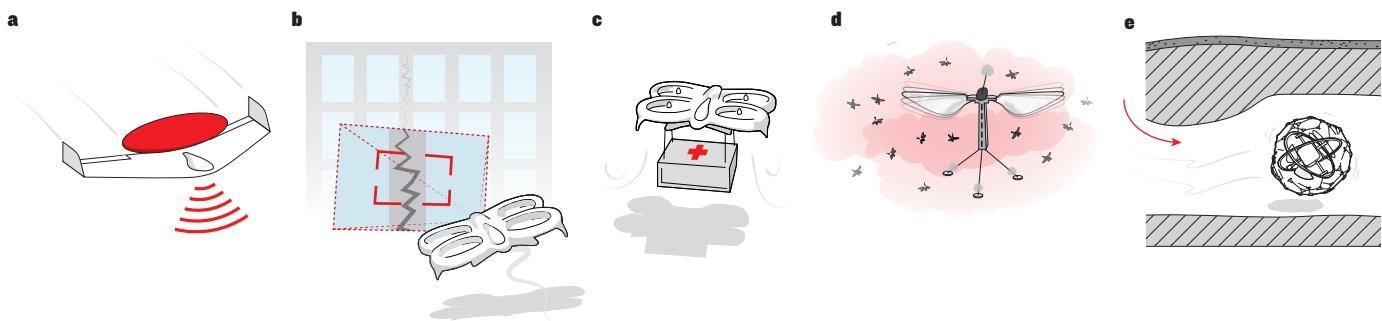


Figure 1 | Autonomous drones in rescue situations. **a**, Fixed-wing drones with a long flight time could provide bird's-eye-view images and a communication network for rescuers on the ground. **b**, Rotorcrafts with hovering capabilities could inspect structures for cracks and leaks; and **c**, transport medical supplies from nearby hospitals. **d**, Swarms of dispensable drones with flapping wings could enter buildings to search for chemical hazards. **e**, Multi-modal caged robots could fly and roll into complex structures to safely search for signs of life.

craft nor rotorcraft scale down well — both in terms of the aerodynamics that govern flight and in the performance of the components that are necessary to generate propulsion. Bio-inspired flapping-wing flight (Fig. 2c) offers an alternative paradigm that can be scaled down in size, but brings fluid-mechanics-modelling and control challenges.

Propulsion and manoeuvrability considerations

Small-scale aerial vehicles are predominantly fixed-wing propeller-driven aircraft or rotorcraft. The latter can be further broken down into conventional primary and tail rotor configurations, coaxial dual rotors¹, or quadcopter designs, which have recently become popular. For each of these configurations, and for propeller-driven aircraft at larger scales, there have been only small incremental improvements to propeller design over the past several decades, since the advent of standards and tools to generate airfoil shapes by the National Advisory Committee for Aeronautics². Propulsive efficiencies for rotorcraft degrade as the vehicle size is reduced; an indicator of the energetic challenges for flight at small scales. Smaller size typically implies lower Reynolds numbers, which in turn suggests an increased dominance of viscous forces, causing greater drag coefficients and reduced lift coefficients compared with larger aircraft. To put this into perspective, this means that a scaled-down fixed-wing aircraft would be subject to a lower lift-to-drag ratio and thereby require greater relative forward velocity to maintain flight, with the associated drag and power penalty reducing the overall energetic efficiency. The impacts of scaling challenges (Fig. 3) are that smaller drones have less endurance, and that the overall flight times range from tens of seconds to tens of minutes — unfavourable compared with human-scale vehicles.

There are, however, manoeuvrability benefits that arise from decreased vehicle size. For example, the moment of inertia is a strong function of the vehicle's characteristic dimension — a measure of a critical length of the vehicle, such as the chord length of a wing or length of a propeller in a similar manner as used in Reynolds number scaling. Because the moment of inertia of the vehicle scales with the characteristic dimension, L , raised to the fifth power, a decrease in size from a 11 m wingspan, four-seat aircraft such as the Cessna 172 to a 0.05 m rotor-to-rotor separation Blade Pico QX quadcopter implies that the Cessna has about 5×10^{11} the inertia of the quadcopter (with respect to roll). Depending on the scaling law used for torque generated by the quadcopter, this leads to angular acceleration that scales either by L^{-1} or L^{-2} , both cases indicating that the quadcopter will be much more manoeuvrable³. This has resulted in remarkable demonstrations of aerobic manoeuvres in rooms equipped with real-time tracking and control by external computers, including flying through windows at high speeds⁴ and manipulating objects in flight⁵. This enhanced agility, often achieved at the expense of open-loop stability, requires increased emphasis on control — a challenge also exacerbated by the size, weight and power constraints of these small vehicles. Implementing these behaviours in autonomous drones (with on-board sensing and computation) is a significant challenge, but one that is beginning to have promising results^{6,7} — even for aggressive flight manoeuvres. The

challenges for further scaling down these sensing and control systems are discussed in the final section of this Review.

Actuation, power and manufacturing

Beyond autonomy considerations, the size and type of the flying robot also engender associated challenges with actuation and manufacturing. For actuation, rotorcraft and propeller-driven fixed-wing drones generally use electromagnetic motors. Some flapping-wing aircraft^{8,9} also use electromagnetic motors, but require a linkage mechanism to convert the rotary motion of the motor to the flapping motion of the wings. However, as the scale is reduced, conventional motors become inefficient, require substantial gearing to achieve a desired wing or propeller velocity, and are extremely difficult to manufacture. Therefore, for vehicles below a few grams, alternative methods of actuation are required. Similarly, 'macro size' flying robots (tens to hundreds of grams and larger) may be constructed using conventional methods such as additive and subtractive machining and 'nuts-and-bolts' assembly. However, for 'micro size' flying robots (less than a few grams), manufacturing needs to be rethought and novel methods applied. To give a perspective on the challenges for actuation, we need to consider how the physics of scaling¹⁰ affects the performance of electromagnetic motors. At macroscopic scales, motors used in electric cars can achieve an efficiency of nearly 90% with a power density of several kilowatts per kilogram whereas at millimetre scales, existing motors produce a few tens of watts per kilogram of power at less than 50% efficiency¹¹. This reduction in performance and the energetic expense of flight at small scales means that both power source and power distribution are important considerations¹². For vehicles for which the primary propulsion system is also responsible for generating control torques, this does not apply (for example, quadcopters). However, for fixed-wing aircraft and some flapping-wing vehicles, it is essential that most of the power budget must be allocated to staying aloft with a proportionately smaller power for actuating control surfaces. For example, nearly 90% of the power budget of the RoboBee (Fig. 2c) is dedicated to the primary power actuators that generate lift to stay in the air¹³. As already discussed, for power actuators, the typical choice is electromagnetic motors. There are alternatives, including electroactive materials such as piezoelectric actuators¹⁴ or ultrasonic motors¹⁵, however, these are most appropriate for the smaller end of the size spectrum. Control actuators have more options at most scales, including voice coils, shape memory alloys¹⁶ and electroactive polymers¹⁷.

Given the challenges for small-scale propulsion, an unsteady approach to force production is potentially a viable option for small vehicles. Unlike fixed-wing aircraft and rotorcraft that produce lift and thrust in a quasi-steady way (either directly in the case of rotorcraft or indirectly by the movement of air over the wings of a fixed-wing drone), the aerodynamics of flapping-wing drones is complex and involves the generation and manipulation of vortical structures in the air¹⁸. As is obvious from its existence in nature, flapping-wing propulsion is an effective form of locomotion, but translating this into

functional designs is not a trivial exercise given the morphological diversity of the flight apparatus of insects, bats and birds. Nonetheless, drone researchers have effectively reproduced natural flight in many different ways. Examples include robotic insects such as RoboBee¹⁹ and the four-winged DelFly⁸, and bird-sized drones such as the Nano Hummingbird⁹ (Fig. 2c).

As size is reduced, in addition to the challenges for actuator selection and manufacturing, there are also considerations of how to manufacture the entire vehicle. At scales for which electromagnetic actuation and bearing-based rotary joints are feasible, more conventional manufacturing methods are used, such as subtractive machining, additive printing and moulding of composite materials. At small scales — for example, for aircraft of comparable size to insects and small birds — some of these techniques fail, typically owing to limited resolution. Alternative methods have been developed; for example, those based on folding (an inherently scalable technique) have been used to create insect-sized robots, avoiding the challenges that are inherent to macro-scale nuts-and-bolts approaches²⁰.

BOX 1

Control autonomy

According to the definition by the International Organization for Standardization, robot autonomy is the ability to perform intended tasks based on current state and sensing, without human intervention⁹⁸. This definition encompasses a wide range of situations, which demand different levels of autonomy depending on the type of robot and the intended use. For example, although autonomy in tethered robots does not concern energy management, mobile robots with long-range travel may require the capability to decide when to abort the current mission and locate a recharging station. In the case of the small drones discussed here, we can identify three levels of increasing autonomy (Table 1).

- Sensory-motor autonomy: translate high-level human commands (such as to reach a given altitude, perform circular trajectory, move to global positioning system (GPS) coordinates or maintain position) into combinations of platform-dependent control signals (such as pitch, roll, yaw angles or speed); follow pre-programmed trajectory using GPS waypoints.
- Reactive autonomy (requires sensory-motor autonomy): maintain current position or trajectory in the presence of external perturbations, such as wind or electro-mechanical failure; avoid obstacles; maintain a safe or predefined distance from ground; coordinate with moving objects, including other drones; take off and land.
- Cognitive autonomy (requires reactive autonomy): perform simultaneous localization and mapping; resolve conflicting information; plan (for battery recharge for example); recognize objects or persons; learn.

Table 1 | Levels of autonomy: requirements, availability and readiness for market

	Exteroceptive sensors	Computational load	Supervision required	Readiness level	Validated on drone type
Sensory-motor autonomy	None or few	Little	Yes	Deployed	All types
Reactive autonomy	Few and sparse	Medium	Little	Partly deployed	Fixed wing, rotorcraft and flapping wing
Cognitive autonomy	Several and high density	High	None	Not yet deployed	Mostly rotorcraft

Multi-modal drones

In many situations, such as search and rescue, parcel delivery in confined spaces and environmental monitoring, it may be advantageous to combine aerial and terrestrial capabilities. Perching mechanisms could allow drones to land on walls²¹ and power lines²² in order to monitor the environment from a high vantage point while saving energy. Agile drones could move on the ground by using legs in conjunction with retractable²³ or flapping wings²⁴. In an effort to minimize the total cost of transport²⁵, which will be increased by the additional locomotion mode, these future drones may benefit from using the same actuation system for flight control and ground locomotion. The different speed and torque requirements of these two locomotion modes could be reconciled by adapting the wing morphology to the specific situation²⁶, similar to the way vampire bats (*Desmodus rotundus*) use their powerful front limbs when flying or walking²⁷. Alternatively, multi-modal locomotion could be obtained by adding large wheels to the sides of hovering drones²⁸, by embedding the propulsion system in a rolling cage²⁹, or by completely decoupling the spherical cage from the inner rotors by means of a gimbal system³⁰ (Fig. 2b); the latter design allows the drone not only to roll on the ground in any direction, but, because the rotors are protected by a cage, also to safely collide with obstacles or humans without attitude perturbations. Conversely, wings could be added to ground-based robots travelling on rough terrain to extend their jumping distances, stabilize the landing phase and reduce the impact with the ground³¹. In this case as well, the total cost of transport could be reduced by sharing the same actuation system between jumping and wing-deployment mechanisms³². Alternatively, one could use pivoting wings, which minimize drag at take-off, improve the transition from ballistic jump to gliding and maximize the gliding ratio³³. In the future, we may also deploy drones in semi-aquatic environments by using design principles inspired by aquatic birds, such as folding wings for plunge diving or hydrophobic surfaces for dry flight, or by flying squid, which use water-jet propulsion for take off³⁴.

Sensing and control

Conventional unmanned drones that fly at high altitudes regulate their attitude (roll, pitch and yaw) and their position (*x*, *y* and *z*) (Fig. 2d) by continuously monitoring and merging data from an inertial measurement unit (IMU), which contains three-axis accelerometers and gyroscopes, and from a global positioning system (GPS). This technology, which is necessary for sensory-motor autonomy, is now also available in small drones for recreational or professional use. However, GPS information is not sufficiently precise for altitude regulation when flying a few metres above the ground and is not always available or reliable in confined areas, such as cities, forests and buildings. Even if reliable GPS information were available, it would need to be combined with a precise map of the drone's surroundings in order to identify obstacle-free trajectories. However, digital maps, such as those used in car navigation systems, are restricted to traversable roads, do not include three-dimensional information (the height of natural structures, buildings, and bridges, and the presence of cables, poles, and so on) and are not refreshed frequently enough to capture landscape modification. Therefore, small autonomous drones flying at low altitude will need more complex levels of control autonomy and additional sensors to detect distances from the surrounding environment and perform safe and stable trajectories. Vision is a promising sensor modality for small drones because compared with other distance sensors such as sonar, infrared and laser range finders used in terrestrial vehicles, it does not require energy to interrogate the environment, and for comparable mass it can gather richer information and span wider fields of view.

Reactive autonomy

Most efforts in small autonomous drones have focused on achieving reactive autonomy (Box 1) by translating decades worth of neuroethological research on vision-based insect flight into simple control algorithms and lightweight sensors³⁵, with many also serving as validation of biological models. Insect vision relies on compound eyes, which are dense arrays of facets pointing in different directions and spanning large fields of view³⁶.

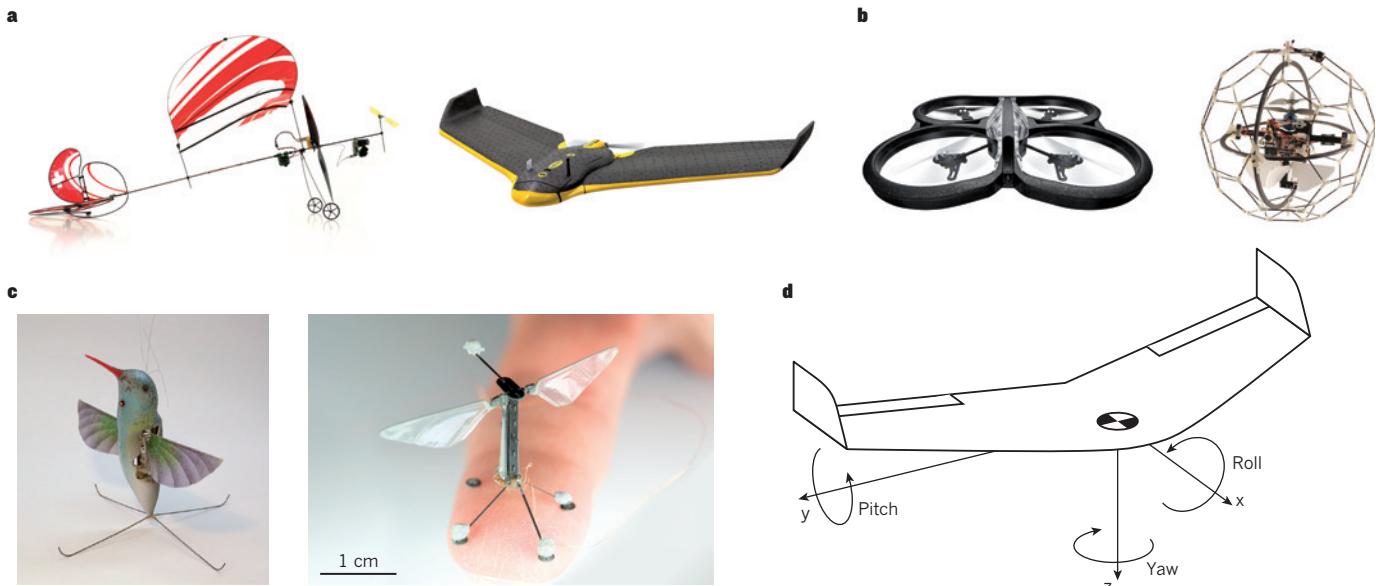


Figure 2 | Drone types with examples. **a**, Fixed-wing drones. A 10 g robot with rudder and flap on the tail is equipped with insect-inspired cameras to avoid obstacles and regulate altitude (left); the 690 g eBee (right) with elevons on the trailing edge of the wings, is equipped with high-resolution cameras for imaging and ventral optic flow sensor for landing. **b**, Rotorcraft. The 380 g AR.Drone 2.0 (left) is a quadcopter equipped with a high-resolution camera for imaging and a ventral optic flow sensor for maintaining position; the 380 g Gimball robot (right) is composed of a

coaxial dual propeller core protected by a decoupled and freely rotating spherical cage. **c**, Flapping-wing drones. The 19 g Nano Hummingbird (left) has a camera for live video streaming; the 80 mg RoboBee (right) can fly tethered to an external power source. **d**, Coordinate system of a generic drone. The drone position is defined in the x , y and z coordinates and the attitude in the yaw, roll and pitch angles. The control algorithms to maintain position and attitude can vary according to drone type, configuration of the actuators and size.

Compound eyes have fixed focus, lower resolution and smaller binocular overlap than human eyes and therefore do not use stereo vision for distance estimation, except for very short ranges³⁷. To safely fly in cluttered environments, insects instead rely on image motion, also known as optic flow^{38,39}, generated by their own displacement relative to the surroundings⁴⁰. It has been experimentally shown that their neural system reacts to optic flow patterns^{41,42} to produce a large variety of flight capabilities, such as obstacle avoidance^{40,43}, speed maintenance⁴⁴, odometry estimation⁴⁵, wall following and corridor centring⁴⁶, altitude regulation^{47,48}, orientation control⁴⁹ and landing^{50,51}. Optic flow intensity is proportional to the distance from objects only during translational movements, but not during rotational movements when it is proportional to the rotational velocity of the agent. Furthermore, optic flow intensities also depend on the speed-to-distance ratio, which raises the issue of how insects can estimate ground speed and distance from obstacles at the same time⁵².

Many vision-based insect capabilities have been replicated with small drones. For example, it has been shown that small fixed-wing drones⁵³ and helicopters⁵⁴ can regulate their distance from the ground using ventral optic flow while a GPS was used to maintain constant speed and an IMU was used to regulate roll angle. The addition of lateral optic flow sensors also allowed a fixed-wing drone to detect near-ground obstacles⁵⁵. Optic flow has also been used to perform both collision-free navigation and altitude control of indoor⁵⁶ and outdoor⁵⁷ fixed-wing drones without a GPS. In these drones, the roll angle was regulated by optic flow in the horizontal direction and the pitch angle was regulated by optic flow in the vertical direction, while the ground speed was measured and maintained by wind-speed sensors. In this case, the rotational optic flow was minimized by flying along straight lines interrupted by short turns or was estimated with on-board gyroscopes and subtracted from the total optic flow, as suggested by biological models^{58,59}. Other authors have even proposed bio-inspired control methods that do not require absolute speed measurement with dedicated sensors, thus eliminating a potential source of error. These methods consist of continuously adjusting flight speed and altitude to maintain constant optic flow signals, which has also been suggested by biological models⁶⁰. For example, altitude control and landing

was achieved by adding negative feedback from ventral optic flow, either to the control surfaces that regulate pitch angle⁵³ or to those that regulate thrust⁶¹. The latter method has also been shown to be effective for altitude control and landing on mobile platforms⁶² and, when used in conjunction with lateral optic flow and lateral thrust, also for replicating flight trajectories of honeybees during indoor flight⁶². However, so far this method has been validated only on tethered drones. An error correction method has also been demonstrated on a quadcopter equipped with omnidirectional vision to fly in corridors by continuously making corrections aimed at reducing the difference between measured optic flow and optic flow expected from flight at the desired altitude, attitude, speed, heading and distance from walls⁶³. Quadcopter and flapping-wing drones are intrinsically unstable and must compensate for positional drift generated by noise in gyroscope and accelerometer signals to hover in place and maintain attitude. It has been shown that the direction of optic flow can be used to reduce uncertainty in inertial sensor signals⁶⁴. Wide-field optical sensors inspired by the simple eyes of insects, called ocelli⁶⁵, have been used for attitude stabilization of flapping-wing drones⁶⁶ and quadcopters⁶⁷.

The need to detect optic flow over wide fields of view at high temporal frequency in a small package has also driven the development of insect-inspired vision sensors that are smaller, lighter and faster than single-lens conventional cameras. For example, some authors^{55–57} have resorted to using the multiple optic flow sensors found in computer optical mice. Whereas others^{68,69} have developed neuromorphic chips that not only extract optic flow, but also adapt to the large variety of light intensities that can be experienced when flying in confined spaces. Similarly, specialized micro optic flow sensors have been used for altitude regulation in sub-1 g flying robots⁷⁰. Miniature, curved, artificial compound eyes with a high density of photoreceptors, wide fields of view and computational properties similar to insect eyes have been recently described^{71,72}, but have not yet been used in drones.

Coordinated flight of multiple drones is another instance of reactive autonomy, but raises additional challenges in sensing, communication and control⁷³. Although very little is known about the sensing and control that underlies coordinated flight in insects, it has been shown that

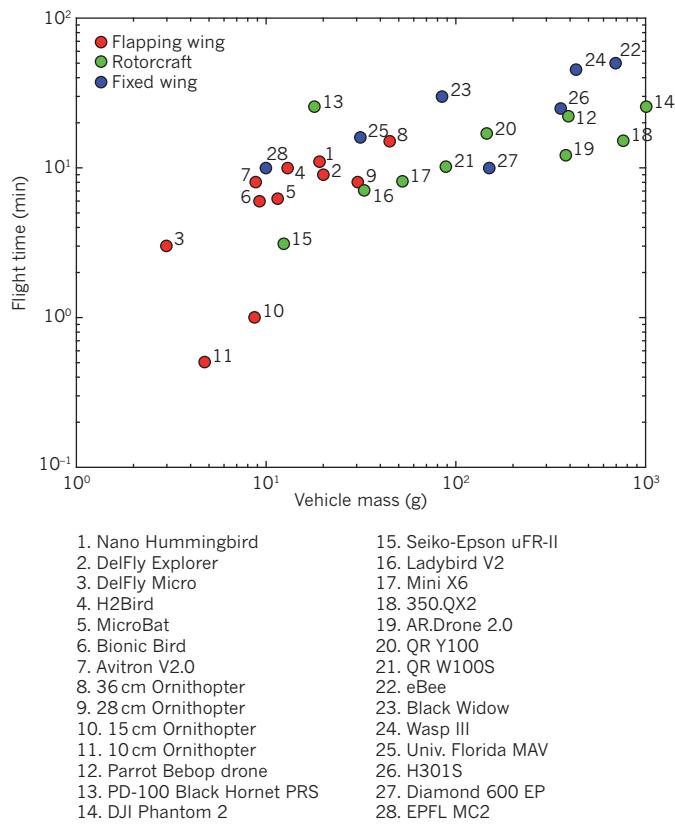


Figure 3 | Flight time against mass of small (less than 1 kg) drones.

Examples include each of the drone types shown in Fig. 2 (fixed wing, rotary and flapping wing). Regardless of the type, there is a clear trend in how flight time scales with mass. Smaller drones have significantly reduced flight times (tens of seconds compared with tens of minutes for larger drones). This is due to actuation limitations and the physics of flight at small scales, as discussed in this Review, and brings about challenges for all levels of autonomy described in Box 1.

cohesion in flocks of starlings is the result of individual birds aligning their trajectories with a small number of their nearest neighbours defined in topological — not metric — space⁷⁴. Evidence from pigeon flocks also suggests that in some conditions, such as homing, a hierarchical organization emerges whereby the trajectories of more experienced birds are tracked and copied by other birds⁷⁵. Computational models of bird flocking build on the combination of three simple behavioural rules, namely repulsion from near individuals, movement towards far neighbours and alignment with average velocities of neighbours⁷⁶. Remarkably, these rules can also account for hierarchical flocking without explicit signalling of experienced agents⁷⁷. However, perceiving the position of other flying drones is still very challenging. Infrared-based bidirectional communication has been used for detecting the range and bearing of neighbouring drones in indoor environments⁷⁸, but the relatively short range of a few metres prevents its use in outdoor environments. Consequently, recent demonstrations of outdoor flocking of ten fixed-wing drones⁷⁹ and ten hovering drones⁸⁰ relied on radio communication of GPS coordinates between neighbouring drones. An alternative approach to perception of neighbours consists of exploiting the direction and intensity of the sound emitted by neighbouring drone engines⁸¹, which may complement or replace visual cues at night or in foggy conditions.

Cognitive autonomy

A different approach is pursued to endow small drones with cognitive autonomy (Box 1). This involves leveraging decades of research in statistical methods for vision-based navigation of terrestrial vehicles, such as visual odometry^{82,83} and simultaneous localization and mapping⁸⁴. For example, monocular vision has been successfully used for simultaneous

localization and mapping of quadcopters in indoor⁸⁵ and outdoor environments⁸⁶. However, these drones could not avoid obstacles that had not been previously detected and mapped. Stereo vision can provide depth information and has been used in a 4 g flapping-wing vehicle to reactively avoid obstacles⁸⁷ and in larger quadcopters to perform simultaneous localization and mapping⁸⁸. Furthermore, both simultaneous localization and mapping and altitude control have been demonstrated in quadcopters with two sets of stereo cameras, one set pointing forwards and one set pointing downwards⁸⁹. Simultaneous localization and mapping algorithms have also been combined with optic flow methods to estimate distances from the surrounding environment and stabilize the drone⁹⁰.

The statistical approach to cognitive autonomy builds on high-resolution digital cameras that can provide dense clouds of data points and on computationally expensive algorithms for reducing uncertainties. Consequently, cognitive autonomy requires heavier sensory payloads and more powerful computational units, which may explain why, so far, progress has been slower for small drones. Although various forms of reactive autonomy have been demonstrated in several types of drones with a wide range of mass and flight endurance (Fig. 3), simultaneous localization, mapping, and path planning have so far been demonstrated mostly in hovering platforms with a relatively large total mass and restricted flight endurance.

Regulatory issues

Unleashing the socio-economic potential of small drones will require not only translating the scientific and technological advances described above into reliable products, but also creating a regulatory framework that will enable public and private entities to operate in full respect of safety, privacy and security concerns. The rapid development of commercial drones and potential benefits prompted the US Federal Aviation Administration (FAA) to define, in 2013, a road map for the gradual integration in the national airspace system of civil unmanned drones that can fly beyond the operator's line of sight by 2028 (ref. 91). In a first “accommodation” stage, which has already started, the FAA allows operation of unmanned aerial vehicles for specific services by issuing airworthiness certificates on request and by defining a set of standards and procedures. In the second “integration” stage, the FAA aims to implement the standards and procedures for operation of civil unmanned aircrafts while retaining on-demand certification for cases that do not fit in the civil unmanned category, and to establish six test sites with diverse geographical and climate conditions to study safety and privacy issues. In the final “evolution” stage, the FAA aims to continuously refine and update regulations, policies and standards on the basis of evolving technological developments and experiences. The FAA has identified sense-and-avoid and communication-and-control as research priorities for full integration of unmanned drones, but for the foreseeable future it will continue to require continuous supervision by a human operator with certified training. In the European Union the legislation is more fragmented because aircraft below 150 kg are separately regulated by individual member states. Some countries, including the Czech Republic, France, Ireland, Italy, Sweden, Switzerland and the United Kingdom, have national regulations for unmanned aircraft; although these regulations differ between countries. For example, the United Kingdom allows operation of drones of less than 20 kg within the line of sight over congested areas on request to the UK Civil Aviation Authority⁹²; France allows remotely operated aircraft to operate beyond the line of sight and near cities under specific conditions⁹³; and Switzerland allows autonomous operation of drones within the line of sight as long as the operator can regain control⁹⁴ if, for example, the drone loses altitude or is about to collide with something. In 2013, the European Commission defined a road map⁹⁵ for the integration of remotely operated aircraft in the European aviation system using a staged approach, the aims, timeline and research priorities of which are similar to those defined in the US road map. However, the first stage of the European road map consists of extending European regulations for large unmanned aircraft systems to those below 150 kg by adapting to national regulations, which may result in more liberal regulations than those in the United States. Furthermore,

the European Commission explicitly aims to improve the public perception of drones and to promote widespread use of these technologies for public and commercial use⁹⁶.

The future of small drones

Scaling up the use of small drones from a niche market to widespread use in civilian applications depends on two related prerequisites: the capability to autonomously and safely manoeuvre in confined spaces and the removal of the legal requirement of supervised operation within the line of sight. We do not foresee major scientific or technological roadblocks to achieving higher levels of autonomous control in research and commercial drones within the next five years. However, the legal requirement of a certified human operator within the line of sight of every single drone is a roadblock that will almost certainly stay for the next five years in the United States and Europe, and removing it will depend on the reliability and safety of small drones.

Assuming that the legal roadblock will gradually be lifted, we expect that reactive forms of control autonomy (Box 1) will become widely available within the next 5–10 years for small commercial drones for long-range operation. We also anticipate that bio-inspired approaches will dominate because they require relatively simple computation and sensors. For example, some commercial drones (such as the eBee and the AR.Drone 2.0 in Fig. 2) already use ventral optic flow for outdoor landing and for indoor position stabilization. In this context, an interesting scientific challenge will be to understand how different navigation capabilities can be integrated into a coherent control system akin to the nervous system of a flying insect. In parallel, an important engineering challenge will be to define test conditions and performance standards in cooperation with governmental institutions and industrial associations to assess the capabilities and reliability of drone technologies.

We also expect rapid progress in cognitive autonomy, which will continue to be driven by the development of artificial intelligence for smartphones capable of identifying human users, learning their behaviours and creating representations of their environment (<http://www.google.com/atap>). On the one hand, face recognition and gesture-based interaction without wearable devices will become widely available for hobby and toy drones within the next five years; for example, by equipping small drones with human-motion-sensing devices developed by the gaming industry. On the other hand, mapping and path planning for autonomous flight in partially unknown and changing environments will represent a challenge for small drones for at least the next ten years and will continue to be illegal in the United States and Europe until at least 2028.

Despite these legal roadblocks, we expect an increasing demand for small drones in civilian use because of their intrinsic safety. Kinetic energy — one measure of the potential of a drone to cause physical harm to a human⁹⁷ — is linearly proportional to the drone's mass and is quadratic in velocity. Thus smaller drones are likely to cause proportionally less harm as the size is reduced, but more subtly, small drones typically operate at lower speeds, dramatically decreasing the potential for harm. For example, a 500 g drone flying at 5 m s⁻¹ has 6.5 J of kinetic energy, which is equivalent to the potential energy of a large apple dropped from about 2 m. As the mobile computing industry continues along the path of miniaturization, drone developers will continue to reap the benefits in the form of smaller, lower power sensor packages (for example, the IMUs used in cell phones and video game controllers) and the slow, but steady increase in battery energy density. Another hardware advance that will inevitably affect future drones is the movement away from general-purpose computation in favour of more specialized high-performance and low-power hardware accelerators tuned for the various functions needed by autonomous drones (both reactive and cognitive autonomy). These accelerators are already present in various mobile devices and are suggested to be a solution to computation for the control of insect-scale drones¹⁹. ■

Received 17 November 2014; accepted 18 March 2015.

1. Bouabdallah, S. *Design and Control of Quadrotors with Applications to*

Autonomous Flying. PhD thesis, Ecole Polytechnique Federale de Lausanne (2007).

2. Ladson, C. L., Brooks, C. W. Jr, Hill, A. S. & Sproles, D. W. *Computer Program to Obtain Ordinates for NACA Airfoils*. Report No. 4741 (NASA, 1996).
3. Kumar, V. & Michael, N. Opportunities and challenges with autonomous micro aerial vehicles. *Int. J. Robot. Res.* **31**, 1279–1291 (2012).

This paper describes scaling laws for dynamics and control of quad rotors and related micro air vehicles.

4. Mellinger, D., Michael, N. & Kumar, V. Trajectory generation and control for precise aggressive maneuvers with quadrotors. *Int. J. Robot. Res.* **31**, 664–674 (2012).
5. Brescianini, D., Hehn, M. & D'Andrea, R. Quadcopter pole acrobatics. In *Proc. International Conference on Intelligent Robots and Systems* 3472–3479 (2013).
6. Shen, S., Mulgaonkar, Y., Michael, N. & Kumar, V. Vision-based state estimation and trajectory control towards aggressive flight with a quadrotor. In *Proc. Robotics: Science and Systems* <http://www.roboticsproceedings.org/rss09/p32.html> (2013).
7. Bry, A., Bachrach, A. & Roy, N. State estimation for aggressive flight in GPS-denied environments using onboard sensing. In *Proc. IEEE International Conference on Robotics and Automation* <http://dx.doi.org/10.1109/ICRA.2012.6225295> (2012).
8. de Croon, G. C. H. E. et al. Design, aerodynamics and autonomy of the DelFly. *Bioinspir. Biomim.* **7**, 025003 (2012).
9. Keenon, M., Klingebiel, K. & Won, H. Development of the Nano Hummingbird: a tailless flapping wing micro air vehicle. In *Proc. AIAA Aerospace Sciences Meeting* <http://dx.doi.org/10.2514/6.2012-588> (2012).
10. Trimmer, W. S. N. Microrobots and micromechanical systems. *Sens. Actuators* **19**, 267–287 (1989).
11. Wood, R. J. et al. Progress on 'pico' air vehicles. *Int. J. Robot. Res.* **31**, 1292–1302 (2012).
12. Finio, B. M. & Wood, R. J. Distributed power and control actuation in the thoracic mechanics of a robotic insect. *Bioinspir. Biomim.* **5**, 045006 (2010).
13. Ma, K., Chirattananon, P., Fuller, S. & Wood, R. J. Controlled flight of a biologically inspired, insect-scale robot. *Science* **340**, 603–607 (2013).

This paper details the development of a flying robotic insect and the first controlled flight of this vehicle.

14. Wood, R. J., Steltz, E. & Fearing, R. S. Optimal energy density piezoelectric bending actuators. *Sensors Actuators A: Physical* **119**, 476–488 (2005).
15. Phys.org. *World's Lightest Micro-Flying Robot Built by Epson* <http://phys.org/news860.html> (2004).
16. Kovac, M., Guignard, A., Nicoud, J.-D., Zufferey, J.-C. & Floreano, D. A 1.5g SMA-actuated microlidger looking for the light. In *Proc. International Conference on Robotics and Automation* 367–372 (2007).
17. Shintake, J., Rosset, S., Schubert, B. E., Floreano, D. & Shea, H. A foldable antagonistic actuator. *IEEE/ASME Trans. Mechatron.* <http://dx.doi.org/10.1109/TMECH.2014.2359337> (2014).
18. Dickinson, M. H., Lehmann, F. O. & Sane, S. P. Wing rotation and the aerodynamic basis of insect flight. *Science* **284**, 1954–1960 (1999).
19. Wood, R. J., Nagpal, R. & Wei, G.-Y. Flight of the RoboBees. *Sci. Am.* **308** <http://www.scientificamerican.com/article/robobee-project-building-flying-robots-insect-size/> (2013).
20. Sreetharan, P., Whitney, J. P., Strauss, M. & Wood, R. J. Monolithic fabrication of millimeter-scale machines. *J. Micromech. Microeng.* **22**, 055027 (2012).
21. Daler, L., Klaptocz, A., Briod, A., Sitti, M. & Floreano, D. A perching mechanism for flying robots using a fibre-based adhesive. In *Proc. International Conference on Robotics and Automation* 4418–4423 (2013).
22. Moore, J. & Tedrake, R. Magnetic localization for perching UAVs on powerlines. In *Proc. International Conference on Intelligent Robots and Systems* 2700–2707 (2011).
23. Bachmann, R. J., Boria, F. J., Vaidyanathan, R. & Ifju, P. G. A biologically inspired micro-vehicle capable of aerial and terrestrial locomotion. *Mechanism Mach. Theory* **44**, 513–526 (2009).
24. Peterson, K., Birkmeyer, P., Dudley, R. & Fearing, R. S. A wing-assisted running robot and implications for avian flight evolution. *Bioinspir. Biomim.* **6**, 046008 (2011).
25. Gabrielli, G. & von Kármán, T. What price speed? Specific power required for propulsion of vehicles. *Mech. Eng.* **72**, 775–781 (1950).
26. Daler, L., Mintchev, S., Stefanini, C. & Floreano, D. A bioinspired multi-modal flying and walking robot. *Bioinspir. Biomim.* **10**, 016005 (2015).
27. Riskin, D. K. & Hermanson, J. W. Biomechanics: independent evolution of running in vampire bats. *Nature* **434**, 292 (2005).
28. Itasse, M., Moschetta, J.-M., Ameho, Y. & Carr, R. Equilibrium transition study for a hybrid MAV. *Inter. J. Micro Air Vehicles* **3**, 229–246 (2011).
29. Kalantari, A. & Spenko, M. Design and experimental validation of HyTAQ, a hybrid terrestrial and aerial quadrotor. In *Proc. International Conference on Robotics and Automation* 4445–4450 (2013).
30. Briod, A., Kornatowski, P. M., Zufferey, J.-C. & Floreano, D. A collision resilient flying robot. *J. Field Robot.* **31**, 496–509 (2014).
31. Vidyasagar, A., Zufferey, J.-C., Floreano, D. & Kovac, M. Performance analysis of jump-gliding locomotion for miniature robotics. *Bioinspir. Biomim.* **10**, 025006 (2015).
32. Woodward, M. A. & Sitti, M. MultiMo-Bat: a biologically inspired integrated jumping-gliding robot. *Int. J. Robot. Res.* **33**, 1511–1529 (2014).
33. Desbiens, A. L., Pope, M. T., Christensen, D. L., Hawkes, E. W. & Cutkosky, M. R. Design principles for efficient, repeated jump gliding. *Bioinspir. Biomim.* **9**, 025009 (2014).
34. Siddall, R. & Kovač, M. Launching the AquaMAV: bioinspired design for

aerial-aquatic robotic platforms. *Bioinspir. Biomim.* **9**, 031001 (2014).

35. Floreano, D., Zufferey, J.-C., Srinivasan, M. V. & Ellington, C. *Flying Insects and Robots* (Springer, 2009).
- This book provides an introduction to insect-inspired drones for biologists and engineers.
36. Land, M. F. & Nilsson, D.-E. *Animal Eyes* (Oxford Univ. Press, 2002).
37. Srinivasan, M. V. How insects infer range from visual motion. *Rev. Oculomot. Res.* **5**, 139–156 (1993).
38. Gibson, J. J. *The Perception of the Visual World* (Houghton Mifflin, 1950).
39. Koenderink, J. J. Optic Flow. *Vision Res.* **26**, 161–179 (1986).
40. Lehrer, M., Srinivasan, M. V., Zhang, S. W. & Horridge, G. A. Motion cues provide the bee's visual world with a third dimension. *Nature* **332**, 356–357 (1988).
41. Franceschini, N., Riehle, A. & Le Nestour, A. in *Facets of Vision* (eds Stavenga, D. G. & Hardie, R. C.) 360–390 (Springer, 1989).
42. Krapp, H. G. & Hengstenberg, R. Estimation of self-motion by optic flow processing in single visual interneurons. *Nature* **384**, 463–466 (1996).
43. Tammero, L. F. & Dickinson, M. H. The influence of visual landscape on the free flight behavior of the fruit fly *Drosophila melanogaster*. *J. Exp. Biol.* **205**, 327–343 (2002).
44. Barron, A. & Srinivasan, M. V. Visual regulation of ground speed and headwind compensation in freely flying honey bees (*Apis mellifera* L.). *J. Exp. Biol.* **209**, 978–984 (2006).
45. Srinivasan, M. V., Zhang, S. W., Altwein, M. & Tautz, J. Honeybee navigation: nature and calibration of the "odometer". *Science* **287**, 851–853 (2000).
46. Serres, J., Masson, G., Ruffier, F. & Franceschini, N. A bee in the corridor: centering and wall-following. *Naturwissenschaften* **95**, 1181–1187 (2008).
47. Portelli, G., Ruffier, F. & Franceschini, N. Honeybees change their height to restore optic flow. *J. Comp. Physiol. A Neuroethol. Sens. Neural Behav. Physiol.* **196**, 307–313 (2010).
48. Straw, A. D., Serlin, L. & Dickinson, M. H. Visual control of altitude in flying *Drosophila*. *Curr. Biol.* **20**, 1550–1556 (2010).
49. Egelhaaf, M. & Borst, A. A look into the cockpit of the fly: visual orientation, algorithms, and identified neurons. *J. Neurosci.* **13**, 4563–4574 (1993).
50. Wagner, H. Flow-field variables trigger landing in flies. *Nature* **297**, 147–148 (1982).
51. Baird, E., Boeddeker, N., Ibbotson, M. R. & Srinivasan, M. V. A universal strategy for visually guided landing. *Proc. Natl Acad. Sci. USA* **110**, 18686–18691 (2013).
52. Taylor, G. K. & Krapp, H. G. Sensory systems and flight stability: what do insects measure and why? *Adv. Insect Physiol.* **34**, 231–316 (2007).
53. Chahl, J. S., Srinivasan, M. V. & Zhang, S. W. Landing strategies in honeybees and applications to uninhabited airborne vehicles. *Int. J. Robot. Res.* **23**, 101–110 (2004).
54. Garratt, M. A. & Chahl, J. S. Vision-based terrain following for an unmanned rotorcraft. *J. Field Robot.* **25**, 284–301 (2008).
55. Griffiths, S. et al. Maximizing miniature aerial vehicles. *IEEE Robot. Autom. Mag.* **13**, 34–43 (2006).
56. Zufferey, J.-C., Klaptocz, A., Beyeler, A., Nicoud, J.-D. & Floreano, D. A 10-gram vision-based flying robot. *Adv. Robot.* **21**, 1671–1684 (2007).
57. Beyeler, A., Zufferey, J.-C. & Floreano, D. Vision-based control of near-obstacle flight. *Auton. Robots* **27**, 201–219 (2009).
58. Chan, W. P., Prete, F. & Dickinson, M. H. Visual input to the efferent control system of a fly's "gyroscope". *Science* **280**, 289–292 (1998).
59. Collett, T. S. Some operating rules for the optomotor system of a hoverfly during voluntary flight. *J. Comp. Physiol. A Neuroethol. Sens. Neural Behav. Physiol.* **138**, 271–282 (1980).
60. Baird, E., Srinivasan, M. V., Zhang, S. & Cowling, A. Visual control of flight speed in honeybees. *J. Exp. Biol.* **208**, 3895–3905 (2005).
61. Ruffier, F. & Franceschini, N. Optic flow regulation: the key to aircraft automatic guidance. *Robot. Auton. Syst.* **50**, 177–194 (2005).
62. Roubieu, F. L. et al. A biomimetic vision-based hovercraft accounts for bees' complex behaviour in various corridors. *Bioinspir. Biomim.* **9**, 036003 (2014).
63. Conroy, J., Gremillion, G., Ranganathan, B. & Humbert, S. J. Implementation of wide-field integration of optic flow for autonomous quadrotor navigation. *Auton. Robots* **27**, 189–198 (2009).
64. Briod, A., Zufferey, J.-C. & Floreano, D. Optic-flow based control of a 46 g quadrotor. In *Proc. Workshop on Vision-based Closed-Loop Control and Navigation of Micro Helicopters in GPS-denied Environments* http://rpg.ifki.uzh.ch/IROS13_TOC.html (2013).
65. Schuppe, H. & Hengstenberg, R. Optical properties of the ocelli of *Calliphora erythrocephala* and their role in the dorsal light response. *J. Comp. Physiol. A Neuroethol. Sens. Neural Behav. Physiol.* **173**, 143–149 (1993).
66. Fuller, S. B., Karpelson, M., Censi, A., Ma, K. Y. & Wood, R. J. Controlling free flight of a robotic fly using an onboard vision sensor inspired by insect ocelli. *J. R. Soc. Interface* **11**, 20140281 (2014).
67. Gremillion, G., Humbert, J. S. & Krapp, H. G. Bio-inspired modeling and implementation of the ocelli visual system of flying insects. *Biol. Cybern.* **108**, 735–746 (2014).
68. Ruffier, F. & Franceschini, N. Optic flow regulation in unsteady environments: a tethered MAV achieves terrain following and targeted landing over a moving platform. *J. Intell. Robot. Syst.* <http://dx.doi.org/10.1007/s10846-014-0062-5> (2014).
69. Ruffier, F., Viollet, S., Amic, S. & Franceschini, N. Bio-inspired optical flow circuits for the visual guidance of micro air vehicles. In *Proc. International Symposium on Circuits and Systems* **3**, 846–849 (2003).
70. Duhamel, P.-E. J., Pérez-Arancibia, N. O., Barrows, G. & Wood, R. J. Biologically

inspired optical-flow sensing for altitude control of flapping-wing microrobots. *IEEE Trans. Mechatron.* **18**, 556–568 (2013).

71. Floreano, D. et al. Miniature curved artificial compound eyes. *Proc. Natl Acad. Sci. USA* **110**, 9267–9272 (2013).
72. Song, Y. M. et al. Digital cameras with designs inspired by the arthropod eye. *Nature* **497**, 95–99 (2013).
73. Zufferey, J. C. et al. Aerial collective systems. In *Handbook of Collective Robotics* (ed. Kernbach, S.) 609–660 (CRC, 2013).

This chapter provides an introduction and survey of collective aerial vehicles.

74. Ballerini, M. et al. Interaction ruling animal collective behavior depends on topological rather than metric distance: evidence from a field study. *Proc. Natl Acad. Sci. USA* **105**, 1232–1237 (2008).
75. Nagy, M., Akos, Z., Biro, D. & Vicsek, T. Hierarchical group dynamics in pigeon flocks. *Nature* **464**, 890–893 (2010).
76. Reynolds, C. W. Flocks, herds, and schools: a distributed behavioral model. *Comput. Graph.* **21**, 25–34 (1987).
77. Couzin, I. D., Krause, J., Franks, N. R. & Levin, S. A. Effective leadership and decision-making in animal groups on the move. *Nature* **433**, 513–516 (2005).
78. Roberts, J., Stirling, T., Zufferey, J.-C. & Floreano, D. 3-D relative positioning sensor for indoor flying robots. *Auton. Robots* **33**, 5–20 (2012).
79. Hauer, S. et al. Reynolds flocking in reality with fixed-wing robots: communication range vs. maximum turning rate. In *Proc. International Conference on Robots and Systems* 5015–5020 (2011).
80. Virág, C. et al. Flocking algorithm for autonomous flying robots. *Bioinspir. Biomim.* **9**, 025012 (2014).
81. Basiri, M., Schill, F. S., Floreano, D. & Lima, P. Audio-based localization for swarms of micro air vehicles. In *Proc. International Conference on Robotics and Automation* <http://infoscience.epfl.ch/record/196274> (2014).
82. Scaramuzza, D. & Fraundorfer, F. Visual odometry: part I — the first 30 years and fundamentals. *IEEE Robot. Autom. Mag.* **18**, 80–92 (2011).
83. Fraundorfer, F. & Scaramuzza, D. Visual odometry: part II — matching, robustness, and applications. *IEEE Robot. Autom. Mag.* **19**, 78–90 (2012).
84. Davison, A. J. & Murray, D. W. Simultaneous localization and map-building using active vision. *IEEE Trans. Pattern Anal. Mach. Intell.* **24**, 865–880 (2002).
85. Bachrach, A., He, R. & Roy, N. Autonomous flight in unknown indoor environments. *Inter. J. Micro Air Vehicles* **1**, 217–228 (2009).
86. Scaramuzza, D. et al. Vision-controlled micro flying robots: from system design to autonomous navigation and mapping in GPS-denied environments. *IEEE Robot. Autom. Mag.* **21**, 26–40 (2014).

This paper details the development of a drone capable of outdoor simultaneous localization and mapping without GPS.

87. de Wagter, C., Tijmons, S., Remes, B. D. W. & de Croon, G. C. H. E. Autonomous flight of a 20-gram flapping wing MAV with a 4-gram onboard stereo vision system. In *Proc. International Conference on Intelligent Robots and Systems* 4982–4987 (2014).
88. Heng, L. et al. Autonomous visual mapping and exploration with a micro aerial vehicle. *J. Field Robot.* **31**, 654–675 (2014).
89. Schauwecker, K. & Zell, A. On-board dual-stereo-vision for the navigation of an autonomous MAV. *J. Intell. Robot. Syst.* **74**, 1–16 (2014).
90. Kendoul, F., Fantoni, I. & Nonami, K. Optic flow-based vision system for autonomous 3D localization and control of small aerial vehicles. *Robot. Autom. Syst.* **57**, 591–602 (2009).
91. US Department of Transportation. *Integration of Civil Unmanned Aircraft Systems (UAS) in the National Airspace System (NAS) Roadmap* (US Department of Transportation Federal Aviation Authority, 2013).
92. UK Civil Aviation Authority. *Unmanned Aircraft System Operations in UK Airspace — Guidance*. Report CAP 722 (UK Civil Aviation Authority, 2012).
93. Ministère de l'Énergie, du Développement Durable, des Transports et du Logement. *Relatif à l'utilisation de l'espace aérien par les aéronefs qui circulent sans personne à bord* [in French]. Report No. Texte 9 sur 308, (Ministère de l'Énergie, du Développement Durable, des Transports et du Logement, 2012).
94. Swiss Federal Office of Civil Aviation. *Drones and aircraft models* <http://www.baz.admin.ch/dienstleistungen/02658/index.html?lang=en> (Swiss Federal Office of Civil Aviation, 2014).
95. European RPAS Steering Group. *Roadmap for the Integration of Civil Remotely-Piloted Aircraft Systems into the European Aviation System* <http://ec.europa.eu/enterprise/sectors/aerospace/uas/> (European RPAS Steering Group, 2013).
96. European Commission. *A New Era for Aviation: Opening the Aviation Market to the Civil Use of Remotely Piloted Aircraft Systems in a Safe and Sustainable Manner*. Report COM(2014) 207 (European Commission, 2014).
97. Haddon, D. R. & Whittaker, C. J. Aircraft airworthiness certification standards for civil UAVs. *Aeronaut. J.* **107**, 79–86 (2003).
98. International Organization for Standardization. *Robots and robotic devices — vocabulary* ISO 8373:2012 (International Organization for Standardization, 2012).

Acknowledgements D.F. and R.J.W. thank the Wyss Institute for Biologically Inspired Engineering at Harvard University, where this Review was written. D.F. also thanks the Swiss National Science Foundation through the National Centre of Competence in Research Robotics.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of this paper at go.nature.com/xukalu. Correspondence should be addressed to D.F. (dario.floreano@epfl.ch).

Design, fabrication and control of soft robots

Daniela Rus¹ & Michael T. Tolley²

Conventionally, engineers have employed rigid materials to fabricate precise, predictable robotic systems, which are easily modelled as rigid members connected at discrete joints. Natural systems, however, often match or exceed the performance of robotic systems with deformable bodies. Cephalopods, for example, achieve amazing feats of manipulation and locomotion without a skeleton; even vertebrates such as humans achieve dynamic gaits by storing elastic energy in their compliant bones and soft tissues. Inspired by nature, engineers have begun to explore the design and control of soft-bodied robots composed of compliant materials. This Review discusses recent developments in the emerging field of soft robotics.

Biology has long been a source of inspiration for engineers making ever-more capable machines¹. Softness and body compliance are salient features often exploited by biological systems, which tend to seek simplicity and show reduced complexity in their interactions with their environment². Several of the lessons learned from studying biological systems are now culminating in the definition of a new class of machine that we, and others, refer to as soft robots^{3–6}. Conventional, rigid-bodied robots are used extensively in manufacturing and can be specifically programmed to perform a single task efficiently, but often with limited adaptability. Because they are built of rigid links and joints, they are unsafe for interaction with humans. A common practice is to separate human and robotic work spaces in factories to mitigate safety concerns. The lack of compliance in conventional actuation mechanisms is part of this problem. Soft robots provide an opportunity to bridge the gap between machines and people. In contrast to hard-bodied robots, soft robots have bodies made out of intrinsically soft and/or extensible materials (for example, silicone rubbers) that can deform and absorb much of the energy arising from a collision. These robots have a continuously deformable structure with muscle-like actuation that emulates biological systems and results in a relatively large number of degrees of freedom compared with their hard-bodied counterparts. They (Fig. 1) have the potential to exhibit unprecedented adaptation, sensitivity and agility. Soft robots promise to be able to bend and twist with high curvatures and thus can be used in confined spaces⁷; to deform their bodies in a continuous way and thus achieve motions that emulate biology⁸; to adapt their shape to the environment, employing compliant motion and thus manipulate objects⁹, or move on rough terrain and exhibit resilience¹⁰; or to execute rapid, agile manoeuvres, such as the escape manoeuvre in fish¹¹.

The key challenge for creating soft machines that achieve their full potential is the development of controllable soft bodies using materials that integrate sensors, actuators and computation, and that together enable the body to deliver the desired behaviour. Conventional approaches to robot control assume rigidity in the linkage structure of the robot and are a poor fit for controlling soft bodies, thus soft materials require new algorithms.

What is soft?

'Soft' refers to the body of the robot. Soft materials are the key enablers for creating soft robot bodies. Although Young's modulus is only defined for homogeneous, prismatic bars that are subject to

axial loading and small deformations, it is nonetheless a useful measure of the rigidity of materials used in the fabrication of robotic systems⁵. Materials conventionally used in robotics (for example, metals or hard plastics) have moduli in the order of 10^9 – 10^{12} pascals, whereas natural organisms are often composed of materials (for example, skin or muscle tissue) with moduli in the order of 10^4 – 10^9 Pa (orders of magnitude lower than their engineered counterparts; Fig. 2). We define soft robots as systems that are capable of autonomous behaviour, and that are primarily composed of materials with moduli in the range of that of soft biological materials.

The advantages of using materials with compliance similar to that of soft biological materials include a considerable reduction in the harm that could be inadvertently caused by robotic systems (as has been demonstrated for rigid robots with compliant joints¹²), increasing their potential for interaction with humans. Compliant materials also adapt more readily to various objects, simplifying tasks such as grasping¹³, and can also lead to improved mobility over soft substrates¹⁴.

For the body of a soft robot to achieve its potential, facilities for sensing, actuation, computation, power storage and communication must be embedded in the soft material, resulting in smart materials. In addition, algorithms that drive the body to deliver the desired behaviours are required. These algorithms implement impedance matching to the structure of the body. This tight coupling between body and brain allows us to think about soft-bodied systems as machines with mechanical intelligence, in which the body can be viewed as augmenting the brain with morphological computation^{15,16}. This ability of the body to perform computation simplifies the control algorithms in many situations, blurring the line between the body and the brain. However, soft robots (like soft organisms) require control algorithms, which (at least for the foreseeable future) will run on some sort of computing hardware. Although both body and brain must be considered in concert, the challenges involved are distinct enough that, in this Review, we find it useful to organize them into separate sections.

In the following three sections we review recent developments in the field of soft robotics as they pertain to design and fabrication, computation and control, and systems and applications. We then discuss persistent challenges facing soft robotics, and suggest areas in which we see the greatest potential for societal impact.

¹Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, The Stata Center, Building 32, 32 Vassar Street, Cambridge, Massachusetts 02139, USA.

²Department of Mechanical and Aerospace Engineering, University of California, San Diego, 9500 Gilman Drive, La Jolla, California 92093-0403.

Design and fabrication

A robot is classified as hard or soft on the basis of the compliance of its underlying materials³. Soft robots are capable of continuum deformations, but not all continuum robots are soft. For example, the robotic elephant trunk manipulator¹⁷ is a discrete hyper-redundant continuum robot composed of rigid materials; the articulated catheter robot¹⁸ is an example of a hard continuum robot; OctArm¹⁹ (Fig. 3b) is an example of a semi-soft continuum robot; while the caterpillar robot²⁰ (Fig. 1a) and the rolling belt robot²¹ are examples of soft continuum robots. These soft machines have modular bodies consisting of soft rubber segments, which can be composed serially or in parallel to create complex morphologies. The body of a soft robot may consist of multiple materials with different stiffness properties^{11,22}. A soft robot encases in a soft body all the subsystems of a conventional robot: an actuation system, a perception system, driving electronics and a computation system, with corresponding power sources. Technological advances in soft materials and subsystems compatible with the soft body enable the autonomous function of the soft robot. The rest of this section describes recent progress in developing subsystems for soft robots. With this range of components, design tools are used to create the topology of the robot body along with the placement of its functional components. Given the design road map, the robot is ready to be fabricated.

Actuation

The segments of a soft robot are usually actuated in one of two ways (Fig. 4): variable length tendons (in the form of tension cables²³ or shape-memory alloy actuators²⁴) may be embedded in soft segments, to achieve, for example, robotic octopus arms (Fig. 3f); or pneumatic actuation is used to inflate channels in a soft material and cause a desired deformation. Pneumatic artificial muscles (PAMs), also known as McKibben actuators, are examples of compliant linear soft actuators composed of elastomer tubes in fibre sleeves^{25,26}. Fluidic elastomer actuators (FEAs) are a new type of highly extensible and adaptable, low-power soft actuator. FEAs comprise synthetic elastomer films operated by the expansion of embedded channels under pressure. Once pressurized, the actuator will keep its position with

little or no additional energy consumption. FEAs can be operated pneumatically^{10,27–30} or hydraulically^{31,32}. Given a small number of options for pressurized working fluid generation, and a significant difference between the time constants for the generators and actuators, pressure-regulating components such as regulators and valves are necessary. Regardless of the actuation method, soft actuators are frequently arranged in a biologically inspired agonist–antagonist arrangement (like muscles) to allow bidirectional actuation. Another benefit of this arrangement is that co-contraction of muscle pairs leads to adaptable compliance.

The design and actuation of soft-robotic systems dates back to at least 1992, when a team demonstrated the impressive capabilities of soft, flexible microactuators²⁷ (Fig. 3a). This work demonstrated the approach of pneumatic actuation of robotic elements composed of elastomer. In this approach, a fluid (usually air) is used to inflate channels in the elastomer, while some asymmetry in the design or constituent materials causes the component to actuate (move) in the desired way (Fig. 4). The resulting continuous, adaptive motions can seem surprisingly lifelike. Other groups^{29,33–35} (Fig. 2b, c) used soft lithography techniques adapted from microfluidics, as well as soft composite materials composed of various silicone polymers and elastomers at times embedded with paper or cloth, to design and fabricate pneumatically actuated soft systems. One challenge with these designs based on pneumatic networks (also referred to as pneu-nets) is that the high strains required for actuation can lead to slow actuation rates and rupture failures. A slightly more complex design for pneumatically actuated elastomeric soft robots reduced the material strain required for actuation³⁶, and allowed the untethered walking of a large, soft robot¹⁰. Soft-lithography fabrication approaches typically use a layer of stiffer rubber or elastomer, sometimes with paper, fabric or a plastic film embedded, to achieve asymmetric strain for actuation. An alternative approach is to augment all elastomeric elements with flexible fibres, which limit the stress taken up by the elastomer during pneumatic actuation. The result is soft actuators with reduced extensibility and flexibility, but the ability to withstand higher actuation pressures, and hence apply larger forces. Using complex moulding and/or free-form fabrication techniques,

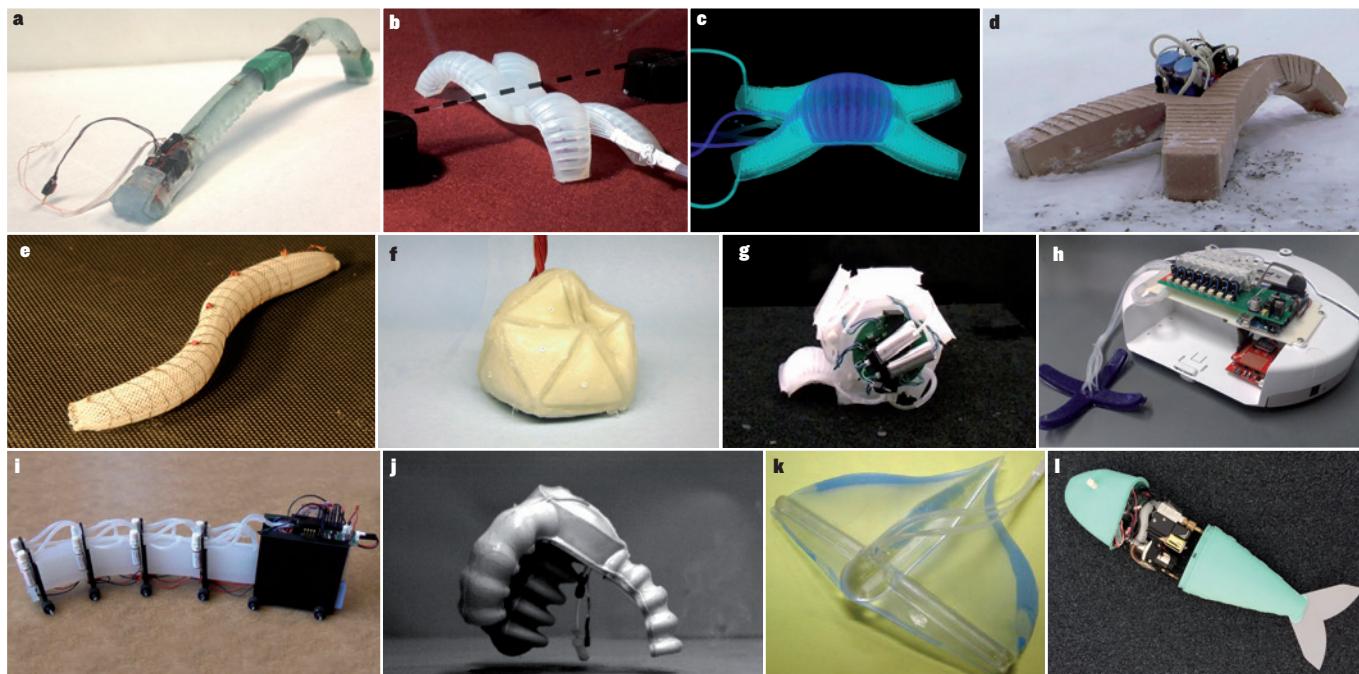


Figure 1 | Mobile soft-robotic systems inspired by a range of biological systems. **a**, Caterpillar-inspired locomotion²⁰. **b**, A multi-gait quadruped²⁹. **c**, Active camouflage³⁵. **d**, Walking in hazardous environments¹⁰. **e**, Worm-inspired locomotion⁸⁸. **f**, Particle-jamming-based actuation⁴². **g**, Rolling powered by a pneumatic battery²⁸. **h**, A hybrid hard–soft robot⁸⁹. **i**, Snake-inspired locomotion⁸. **j**, Jumping powered by internal combustion⁵⁸. **k**, Manta-ray inspired locomotion¹⁰⁰. **l**, An autonomous fish¹¹.

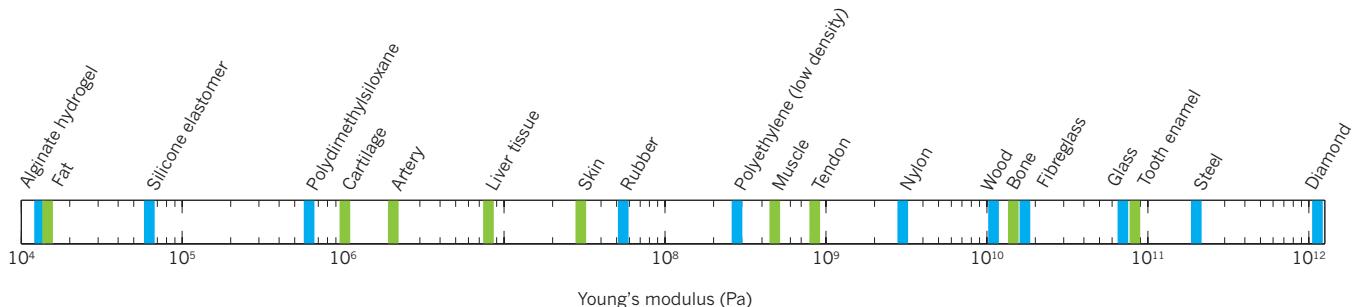


Figure 2 | Approximate tensile modulus (Young's modulus) of selected engineering and biological materials. Soft robots are composed primarily of materials with moduli comparable with those of soft biological materials (muscles, skin, cartilage, and so on), or of less than around 1 gigapascal. These materials exhibit considerable compliance under normal loading conditions.

it is possible to embed fibres directly into pneumatic elastomeric actuators to achieve agile motions based on bending^{13,27,37}.

Although most soft-robot prototypes have used pneumatic or hydraulic actuation, a great deal of research has focused on the development of electrically activated soft actuators composed of electroactive polymers (EAPs)^{38,39}, which have also led to prototype systems²². Since energy is typically most readily stored in electrical form, and computation is usually done on electronic circuits, it may be more efficient to directly use electrical potential to actuate soft robots. Types of EAPs include dielectric EAPs, ferroelectric polymers, electrostrictive graft polymers, liquid crystal polymers, ionic EAPs, electrorheological fluids, ionic polymer–metal composites, and stimuli-responsive gels. Since a detailed discussion is beyond the scope of this Review, we refer the reader to refs 38, 39 for details. In general, fabrication, performance, and long-term stability are active areas of research in EAPs.

Instead of designing the stiffness of robotic systems by tuning their constituent materials, another line of soft-robotics research has sought to control material stiffness on the fly. One approach is to embed or encase soft materials with stiffer materials such as wax⁴⁰ or metal⁴¹, which can be thermally softened. Embedded heaters can thus be used to adjust the structure's effective stiffness and allow for compliant behaviour or actuated repositioning. Similarly, isothermal phase change caused by particle jamming has also been explored as a method of adjusting a soft robot's rigidity for actuation^{42,43} (Fig. 2f), or even for grasping an impressive array of objects⁴⁴ (Fig. 3c).

Stretchable electronics

So far, most integrated soft-robotic systems have relied on conventional, rigid electronics to store the control algorithms and connect to the systems' actuators, sensors and power sources. However, there has recently been much research in the area of soft and stretchable electronics^{45–47}. A full discussion of this area is beyond the scope of this Review, but, as this field of electronics matures, we expect greater integration with soft robots, resulting in completely soft prototypes.

Sensing

The compliance and morphology of soft robots precludes the use of many conventional sensors including encoders, metal or semiconductor strain gauges, or inertial measurement units (IMUs). Although flexible-bending sensors based on piezoelectric polymers are available as commercial products, these may not be appropriate owing to the need for all elements of the system to be both bendable and stretchable. Soft, stretchable electronics may enable new sensing modalities^{48,49}. The basis of proprioceptive sensors for a soft robot is usually either non-contact sensors or very low modulus elastomers combined with liquid-phase materials. Because soft robots are actuated by generating curvatures, proprioception relies on curvature sensors. The low modulus of proposed elastomer sensors (which

have characteristic moduli in the range of 10⁵–10⁶ Pa) impart minimal change on the impedance of the underlying structures. These sensors generally have layered structures, in which multiple thin elastomer layers are patterned with microfluidic channels by soft lithography. The channels are subsequently filled with a liquid conductor (for example, gallium-containing alloys such as eutectic gallium–indium, or EGaIn). With layered channel geometries, it is possible to tailor sensors for measuring various strains including tensile, shear⁵⁰ or curvature⁵¹. To address the fabrication challenges involved in injecting complex channel networks with liquid conductor, recent work has investigated mask deposition of the conductor⁵², or direct 3D printing of conductive material⁵³. Alternatively, exteroceptive sensing may be used to measure the curvature of a soft robot's body segments in real time³⁰. To expand the applications of soft robotics, compatible chemical and biological sensors⁵⁴ may be used to sense environmental signals. Such sensors may be more compatible with soft robots than the optical and audio recorders typically used in robotics.

Power sources

A big challenge for soft robots is stretchable, portable power sources for actuation. For pneumatic actuators, existing fluidic power sources are not soft and are usually big and bulky. Current off-the-shelf pressure sources are generally limited to compressors or pumps, and compressed air cylinders⁵⁵. If we draw an analogy to electrical systems, compressors are similar to generators as they convert electrical energy into mechanical energy, and compressed gas cylinders are similar to capacitors as they store a pressurized fluid in a certain volume to be discharged when required. Miniature compressors use valuable electrical energy inefficiently, and cylinders of useful form factors do not offer longevity. What has been missing for fluidic systems is the equivalent of a battery, whereby a chemical reaction generates the necessary energy for actuation using a fuel. The chemically operated portable pressure source, or pneumatic battery²⁸ (Fig. 2g), generates pressurized gas using a hydrogen peroxide monopropellant⁵⁶. Combustible fuels are another promising high-energy-density chemical fuel source^{57,58}.

Electrically powered actuators (as well as the electrical controllers for pneumatic systems) require soft, flexible, lightweight electrical power sources⁵⁹. As with soft electronics, this is an active area of research. Recent promising developments include batteries based on graphene⁶⁰, organic polymers⁶¹ and embedded conductive fabric⁶².

Design

Existing soft robotic systems have typically been designed with conventional 3D computer-aided design (CAD) software. However, current CAD software was not created with free-form 3D fabrication processes in mind, and does not easily accommodate the complex non-homogeneous 3D designs that may be desired for soft robotics. This has led researchers to either rely on relatively simple '2.5D'

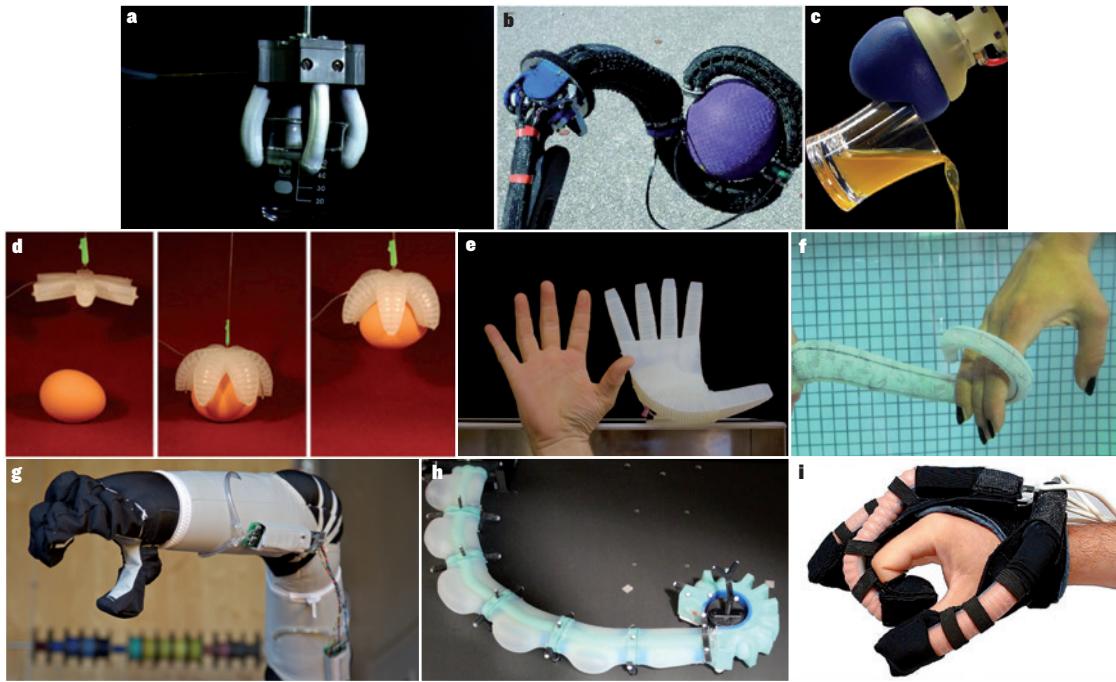


Figure 3 | Grasping and manipulation, which are canonical challenges in robotics, can be greatly simplified with soft robotics. Examples of experimental soft-robotic manipulation systems demonstrating microactuation²⁷ (a), soft-continuum manipulation¹⁹ (b), grasping with

particle jamming⁴⁴ (c), simple gripper fabrication by soft lithography³³ (d), underactuated dexterous grasping¹³ (e), octopus-inspired manipulation²⁴ (f), inflatable robotic manipulators⁹¹ (g), feedback control of a multisegmented arm³⁰ (h) and a soft glove for rehabilitation³² (i).

layered designs, or come up with customized approaches to the design and fabrication of each system, typically based on commercial 3D moulding techniques^{11,58}. Following an alternative approach, researchers have used design automation algorithms inspired by evolution to design soft robots⁶³. Soft-robot designs have been automatically generated using custom finite element analysis software (VoxCAD), which accommodates materials with a large range of moduli, coupled with design optimization using an evolutionary algorithm⁶⁴. In addition, evolutionary algorithms have been used to automatically generate soft-robot designs⁶⁵.

Fabrication

Recent progress in the field of soft robotics has been enabled by the development of rapid digital design and fabrication tools. Researchers have manufactured complex soft-robotic systems by taking advantage of rapid and adaptable fabrication techniques⁶⁶, including multimaterial 3D printing⁶⁷, shape deposition manufacturing (SDM)⁶⁸ and soft lithography⁶⁹. These techniques can be combined to create composites with heterogeneous materials (for example, rubber with different stiffness moduli), embedded electronics and internal channels for actuation^{31,70}. Direct digital printing with soft materials may be another option for fabricating arbitrary soft structures. A variety of robot bodies can be fabricated using these techniques. The next section discusses control systems supporting a wide range of soft-robot applications.

Computation and control

Unlike the control of rigid bodies, the movements of which can be described by six degrees of freedom (three rotations and three translations about the *x*, *y* and *z* axes), the movements of soft bodies cannot be confined to planar motions. Soft materials are elastic and can bend, twist, stretch, compress, buckle, wrinkle and so on. Such motion can be viewed as offering an infinite number of degrees of freedom, which makes the control of soft robots very challenging. Controlling soft robots requires new approaches to modelling, control, dynamics, and high-level planning.

The muscle analogy for soft actuators has driven a number of biologically inspired approaches to modelling and control for soft materials. The octopus arm is the prototypical example of a highly adaptive, soft actuator in nature that has been the source of inspiration for several biomimetic designs. Octopuses form pseudo-joints and use human-like strategies for precise point-to-point movements such as fetching⁷¹. An understanding of the working principles and control of soft organisms (such as the octopus) has led to a model for the control of soft robots⁷². Likewise, the caterpillar has provided an ideal model for mobile soft robots⁷³. The study of these systems for the development and implementation of soft-robotic systems has, in turn, reflected back on our understanding of the mechanics and control of the associated natural systems⁷⁴.

Modelling and kinematics

The kinematics and dynamics of soft-robotic systems are unlike those of conventional, rigid-bodied robots. When composed of a series of actuation elements, these robots approach a continuum behaviour. In theory, the final shape of the robot can be described by a continuous function, and modelling this behaviour requires continuous mathematics. Because soft robots are different from conventional rigid linkage-based systems, researchers have developed new static, dynamic and kinematic models that capture their ability to bend and flex⁷⁵.

Robots made entirely from soft elastomer and powered by fluids do not yet have well-understood models or planning and control algorithms, primarily because their intrinsic deformation is continuous, complex and highly compliant. In addition, soft robots are often under-actuated; they can contain many passive degrees of freedom, and when driven with low-pressure fluids the available input fluid power is unable to compensate for gravitational loading. Designers often model the kinematics of soft robots using a simplifying assumption that leads to the piecewise constant curvature (PCC) model. The PCC model is equivalent to many other modelling approaches⁷⁵. Building on the PCC model, researchers have developed methods to map the actuation space to the configuration

space. These approaches are robot-specific in that they integrate the morphology of the robot and the characteristics of the actuation system. One approach uses Bernoulli–Euler beam mechanics to predict deformation^{3,75,76}; another develops a relationship between the joint variables and the curvature arc parameters⁷⁷, and this applies to high and medium pressure robots; and a third presents models that describe the deformation of robots actuated with low pressures^{11,28,29}. These models are the basis for the forward kinematics of their respective systems. However, the PCC model does not capture all aspects of soft robots and, in an effort to increase the envelope of the model, non-constant curvature models have been developed⁷⁸.

The inverse-kinematics problem, as posed in ref. 75 (computing the curvatures required to place a specified point on the robot body at a desired location), is more challenging. Existing literature has provided several approaches for semi-soft robots^{77,79}. A limitation of existing approaches to solving the inverse-kinematics problem for linear soft bodies (for example, arms) is that currently neither the whole body, nor the pose of the end effector (which may be important for manipulation, sensing, and so on), are not considered in the solution. Autonomous obstacle avoidance and movement through a confined environment are difficult without a computational solution to the inverse-kinematics problem that is aware of the whole body of the robot in space. Real-time, closed-loop curvature controllers are required that drive the bending of the soft pneumatic body segments of the manipulator despite their high compliance and lack of kinematic constraints. One method for closed-loop curvature control of a soft-bodied planar manipulator³⁰ uses the PCC assumption with a cascaded curvature controller. An alternative visual servo control approach for cable-driven soft-robotic manipulators has also been proposed⁸⁰.

The inverse-kinematics algorithm enables task-space planning algorithms to autonomously position their end effector (or some other part of their body) in free space; manoeuvre in confined environments; and grasp and place objects. These task-space algorithms require the planner to consider the entire robot body (the whole arm, for example, not just the end effector for manipulators). Existing algorithms for whole-body control⁸¹ assume rigid body systems and have not been extended to soft-bodied robots. For whole-body control of soft robots, it is difficult to control the configuration of the whole soft body owing to compliance. One computational approach³⁰ to whole-body planning for soft planar manipulators considers both the primary task of advancing the pose of the end effector of the soft robot, and the secondary task of positioning the changing envelope of the whole robot in relation to the environment^{7,30}.

Control

Researchers have used these models to develop new approaches to low-level control, inverse kinematics, dynamic operations and planning for soft-robotic systems. An important aspect of these algorithms is the use of compliance when the robot interacts with its environment. Compliance allows the robot to adapt its shape and function to objects of unknown or uncertain geometry and is the basis for new control and planning algorithms for soft robots in the presence of uncertainty. For example, soft robots can execute pick and place operations without precise positioning or accurate geometric models of the object to be grasped⁴⁴.

Low-level control for soft robots is by pressure control using pressure transducers or volume control using strain sensors. Pressure control accommodates differences in actuator compliance. Volume control is an avenue to configuration control and supports setting a maximum safe displacement limit. Most fluid-powered soft robots use open-loop valve sequencing to control body-segment actuation. Valve sequencing means that a valve is turned on for some duration of time to pressurize the actuator and then turned off to either hold or deflate it. Many soft-robotic systems use this control

approach^{10,21,28,29,34}. Continuously adjustable, variable pressurization using a fluidic drive cylinder has been demonstrated³⁰. Recent work has sought to develop control elements for pneumatic soft robots (for example, valves), which do not require electrical control signals. In this study, passive valves with memory allow the addressable control of many soft-robotic actuators from a single controlled pressure source⁸².

The dynamics of soft robots open the door for new robot capabilities. Much like a child on a swing set, repeatable positioning of a soft tentacle places it outside of its gravity compensation envelope, where the end effector could accomplish tasks. Physical phenomena common to soft robots — including actuation limits, the self-loading effects of gravity, and the high compliance of the manipulator — can be represented as constraints within a trajectory optimization algorithm that operates on a dynamic model of the soft robot⁷. One example dynamic manoeuvre enables soft robots to interact with humans by quickly grabbing objects directly from the hand of a human⁷. Dynamic control of soft robots can be achieved using a new dynamic model for a spatial soft fluidic elastomer manipulator, a method for identifying all unknown system parameters (the soft manipulator, fluidic actuators and continuous drive cylinders or valving), and a planning algorithm that computes locally optimal dynamic manoeuvres through iterative learning control. This approach represents actuation limits, the self-loading effects of gravity, and the high compliance of the manipulator, as constraints within the optimization.

Dynamic models of hard and semi-soft continuum robots provide examples for a variety of control techniques for soft robots^{76,83,84}.

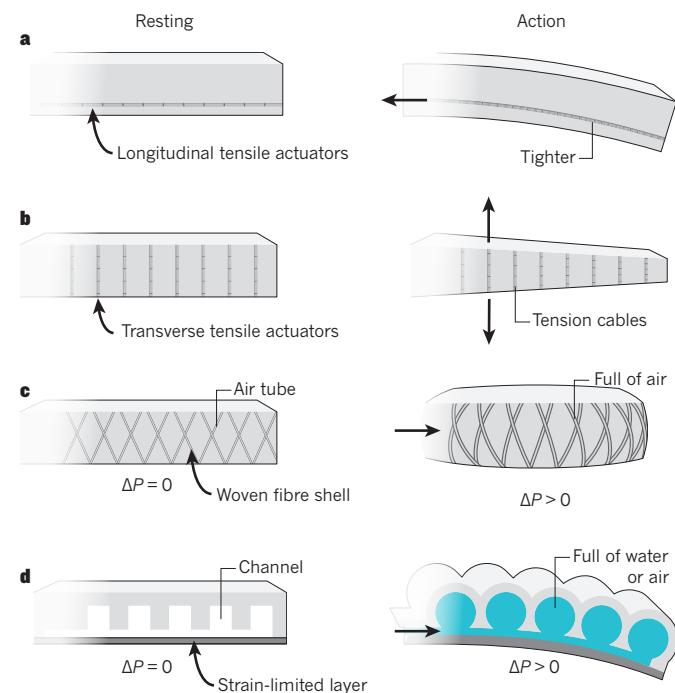


Figure 4 | Cross-section of common approaches to actuation of soft-robot bodies in resting (left) and actuated (right) states. **a**, Longitudinal tensile actuators (for example, tension cables or shape-memory alloy actuators, which contract when heated) along a soft-robot arm cause bending when activated. **b**, Transverse tensile actuators cause a soft-robot arm to extend when contracted (a muscle arrangement also seen in octopuses⁷²). **c**, Pneumatic artificial muscles composed of an elastomeric tube in a woven fibre shell. A pressure applied internally causes the tube and shell to expand radially, causing longitudinal contraction. **d**, Fluidic elastic actuator or PneuNet design consisting of a pneumatic network of channels in an elastomer that expand when filled with a pressurized fluid, causing the soft body to bend toward a strain-limited layer (for example, a stiffer rubber or elastomer embedded with paper or other tensile fibres).

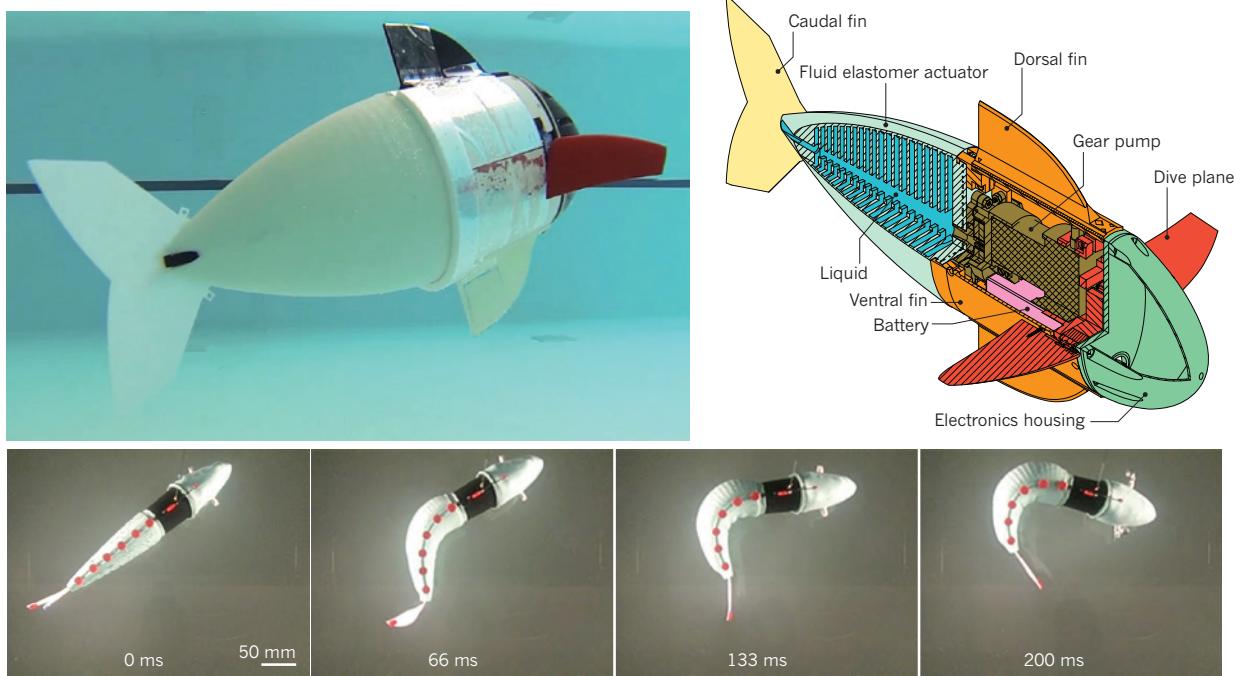


Figure 5 | A soft robotic fish. Soft robotic fish physical prototype (top) and design schematic³¹. Images from a high-speed video of the fish executing a C-turn escape manoeuvre with elapsed time indicated¹¹ (bottom).

Other work has modelled dynamic polymeric pneumatic tubes subject to tip loading using a bending beam approximation⁸⁵, but this has not been used for control. Tatlicioglu *et al.*⁸⁶ developed a dynamic model for, and provided simulations of, a planar extensible continuum manipulator using a Lagrangian approach. Luo *et al.*⁸⁷ modelled the dynamics of a soft planar snake.

Systems and applications

In this section we review some of the systems that have been developed so far to address a variety of potential applications to locomotion, manipulation and human–machine interaction.

Locomotion

Recent work has explored the modes of locomotion possible with (or enabled by) soft bodies (Fig. 1). Notably, studies of caterpillars have led to soft-robotic systems^{20,73}, as well as to a further understanding of the control of motion in these animals⁷⁴. An understanding of worm biomechanics also led to a bioinspired worm design composed of flexible materials and actuated by shape-memory actuators (SMAs)⁸⁸, and an annelid-inspired robot actuated by dielectric elastomer²². A European initiative studied the biomechanics and control of the octopus to produce a soft-robotic prototype^{9,24}. Likewise, a self-contained, autonomous robotic fish actuated by soft fluidic actuators could swim forward, turn and adjust its depth^{11,31} (Fig. 5). This robot can execute a C-turn escape manoeuvre in 100 milliseconds, which is on a par with its biological counterpart, enabling it to be used as an instrument (a physical model) for biological studies. Soft-robotics projects have also explored quadrupedal locomotion^{10,29}, rolling²⁸ and snake-like undulation⁸. Jumping has also been achieved using internal combustion of hydrocarbon fuels to achieve rapid energy release^{57,58}.

Many of the exciting applications for soft robotics (such as search-and-rescue operations or environmental monitoring), require an autonomous, mobile system. However, most experimental soft robotic systems rely on power and/or control signals delivered through pneumatic and/or electrical tethers. Since typical actuation power sources (for example, air compressors or batteries) are

relatively heavy (for example, 1.2 kg¹⁰), tethers greatly simplify system design by significantly reducing the required carrying capacity. One approach to achieving mobile systems is to tether a soft robot to a mobile rigid robot with a greater carrying capacity⁸⁹. Untethered mobile systems have circumvented the challenge of carrying heavy power sources by operating underwater^{11,31} or rolling on a surface^{8,28} such that actuation system masses do not have to be lifted against gravity. Another approach has led to materials and designs tailored to operate at working pressures that are high enough to carry large payloads¹⁰.

Manipulation

Manipulation is one of the canonical challenges for robotics (Fig. 3). Soft systems have a natural advantage over rigid robots in grasping and manipulating unknown objects because the compliance of soft grippers allows them to adapt to a variety of objects with simple control schemes. Grippers that employ isothermal phase change owing to granular jamming have taken advantage of this feature of soft systems^{44,90}. Under-actuated grippers composed of silicone elastomers with embedded pneumatic channels have also shown impressive adaptability^{13,33}. Commercially developed systems have also demonstrated manipulation with lightweight grippers composed of inflated flexible (but not extensible) material⁹¹. As one of the more mature applications of soft-robotic technology, companies have begun to produce soft-robotic manipulators (for example, Pneubotics an Othrlerab company, Empire Robotics and Soft Robotics).

Medical and wearable applications

One of the natural advantages of soft-robotic systems is the compatibility of their moduli with those of natural tissues for medical and wearable applications. Rigid medical devices or orthoses run the risk of causing damage or discomfort to human or animal tissue. In addition, it can be difficult to perfectly replicate the motion of natural joints with rigid orthotics. One possibility is to integrate a degree of compliance into wearable devices, for example for orthopaedic rehabilitation. Recently, researchers have begun to look at medical, wearable applications for soft robotics, including soft wearable

input devices (for example, wearable keyboards⁹²), soft orthotics for ankle–foot rehabilitation³⁷, soft sensing suits for lower-limb measurement⁹³, soft actuated systems for gait rehabilitation of rodents who have had their spinal cord surgically cut⁹⁴, and a soft system for simulation of cardiac actuation⁹⁵.

Soft cyborgs

Recent work has begun to investigate robotic systems that directly integrate biological (as opposed to artificial, biologically compatible) materials. Because biological materials are often very soft, the resulting systems are soft-robotic systems (or perhaps they would be more aptly named soft cyborgs). For example, microbes that digest organic material and produce electricity have powered artificial muscles for autonomous robots⁹⁶, and cardiomyocytes have been used to power a jellyfish-inspired swimming cyborg⁹⁷. One challenge with using swarms of inexpensive soft robots for exploration is retrieving the robots once the task is completed. One way to avoid this problem is to develop biodegradable and soft robots powered by gelatin actuators⁹⁸. Since gelatin is edible, there may also be medical applications for this technology.

Future directions

The field of soft robotics aims to create the science and applications of soft autonomy by asking the questions: how do we design and control soft machines, and how do we use these new machines? Current research on the algorithmic and device-level aspects of soft robots has demonstrated soft devices that can locomote, manipulate, and interact with people and their environment in unique ways. Soft mobile robots capable of locomotion on unknown terrain and soft-robot manipulators capable of pose-invariant and shape-invariant grasping rely on compliance to mitigate uncertainty and adapt to their environment and task. These basic behaviours open the door to applications in which robots work closely with humans. For example, a soft-robot manipulator could handle and prepare a wide variety of food items in the kitchen while a soft mobile robot could use contact to guide patients during physiotherapy exercises. The soft home robot could assist the elderly with monitoring their medicine regimens, whereas a soft factory robot could support humans in delicate assembly tasks. But how do we get to the point where soft robots deliver on their full potential? We need rapid design tools and fabrication recipes for low-cost soft robots, novel algorithmic approaches to the control of soft robots that account for their material properties, tools for developing device-specific programming environments that allow non-experts to use the soft machines, creative individuals to design new solutions and early adopters of the technology.

The soft-robotics community is creating materials that have the ability to compute, sense and move, leading to a convergence between materials and machines. The materials are becoming more like computers owing to embedded electro-mechanical components, and machines are getting softer owing to the alternative materials used to create them. This convergence requires tools for functional specification and automated co-design of the soft body (including all the components necessary for actuation, sensing and computation) and the soft brain (including all software aspects of controlling the device and reasoning about the world and the goal task). Progress in new materials with programmable stiffness properties, proprioceptive sensing, contact modelling, and recipes for rapid fabrication will enable the creation of increasingly more capable soft machines.

But how soft should a soft robot be in order to meet its true potential? As with many aspects of robotics, this depends on the task and environment. For domains that require extreme body compliance (for example, inspecting a pipe with complex geometry or laparoscopic surgery) or dealing with a great amount of uncertainty (for example, locomotion on rocky uneven terrain or grasping unknown objects), soft machines can bring the capabilities of robots to new levels. However, soft robots are difficult to model and control, especially when they need to retain a desired body configuration without external support (for example, an

elephant trunk holding an object at a desired height in the presence of gravity). Augmenting soft machines with skeletons, much like vertebrates, could simplify such tasks. An important consideration for the future of soft machines is how to match the softness of the body to the range of capabilities expected from the machine in a way that most effectively leverages the body.

Rapid fabrication recipes and tools for soft robots with embedded electronics are needed to expand the user base beyond experts. In addition, more advances are needed to expand the computation and control capabilities of existing systems. The current approaches to robot control rely on external localization. But soft robots that operate autonomously need the ability to localize proprioceptively. Current models for soft robots do not capture their dynamics. Improved dynamics models will lead to more capable controllers. Although the current planners for soft robots allow the systems to collide harmlessly with their environments, the collisions are not detected. Contact with the environment, however, can be a useful aspect of task-level planning.

Soft robots have the potential to provide a link between living systems and artificial systems at multiple levels: high-level tasks, in interactions between humans and robots, and in cognition. Pfeifer and Bongard⁹⁹ have argued that bodies are central to the way that we think. This view — embodied artificial intelligence — holds that robotic bodies and brains must be considered and developed in concert, much as the bodies and brains of their biological counterparts co-evolved. Soft-robotic systems have the potential to exploit morphological computation to adapt to, and interact with, the world in a way that is difficult or impossible with rigid systems. Following the principles of embodied artificial intelligence, soft robots may allow us to develop biologically inspired artificial intelligence in ways that are not possible with rigid-bodied robots. ■

Received 27 November 2014, accepted 1 May 2015.

1. Full, R. J. in *Comprehensive Physiology* 853–930 (Wiley, 1997).
2. Dickinson, M. H. et al. How animals move: an integrative view. *Science* **288**, 100–106 (2000).
3. Trivedi, D., Rahn, C. D., Kier, W. M. & Walker, I. D. Soft robotics: biological inspiration, state of the art, and future research. *Appl. Bionics Biomech.* **5**, 99–117 (2008). **This paper discusses the biological inspiration for soft robotics and reviews the first generation of soft-robotic systems, which employ primarily pneumatic artificial muscle or electroactive polymer actuators.**
4. Kim, S., Laschi, C. & Trimmer, B. Soft robotics: a bioinspired evolution in robotics. *Trends Biotechnol.* **31**, 287–294 (2013). **This paper reviews design and actuation approaches for soft robots, and discusses the biomechanics of three organisms that frequently serve as inspiration.**
5. Majidi, C. Soft robotics: a perspective — current trends and prospects for the future. *Soft Robotics* **1**, 5–11 (2014).
6. Laschi, C. & Cianchetti, M. Soft robotics: new perspectives for robot bodyware and control. *Front. Bioeng. Biotechnol.* **2**, 3 (2014).
7. Marchese, A. D., Tedrake, R. & Rus, D. Dynamics and trajectory optimization for a soft spatial fluidic elastomer manipulator. *J. Robotics Res.* (in the press). **This paper describes a new algorithm for dynamic control of soft robots and demonstrates the use of this method to achieve reachability outside of the static envelope of the robot.**
8. Onal, C. D. & Rus, D. Autonomous undulatory serpentine locomotion utilizing body dynamics of a fluidic soft robot. *Bioinspir. Biomim.* **8**, 026003 (2013). **This paper presents a mobile soft robot composed of bidirectional fluidic elastomer actuators to achieve snake-like locomotion.**
9. Mazzolai, B., Margheri, L., Cianchetti, M., Dario, P. & Laschi, C. Soft-robotic arm inspired by the octopus: from artificial requirements to innovative technological solutions. *Bioinspir. Biomim.* **7**, 025005 (2012). **This paper presents the design of a biologically inspired artificial muscular hydrostat for underwater soft robotics.**
10. Tolley, M. T. et al. A resilient, untethered soft robot. *Soft Robotics* **1**, 213–223 (2014). **This paper addresses the design of a soft-robot system with legs that is capable of carrying the components required for untethered walking without a skeleton, and demonstrates that soft robots can be resilient to harsh conditions such as large external forces and extreme temperatures.**
11. Marchese, A. D., Onal, C. D. & Rus, D. Autonomous soft robotic fish capable of escape maneuvers using fluidic elastomer actuators. *Soft Robotics* **1**, 75–87 (2014). **This paper describes a self-contained autonomous robot fish that is capable of executing biologically inspired escape manoeuvres (rapid changes in direction) in hundreds of milliseconds, which is on a par with similarly sized fish.**

12. Albu-Schaffer, A. et al. Soft robotics. *IEEE Robot. Automat. Mag.* **15**, 20–30 (2008).

13. Deimel, R. & Brock, O. A novel type of compliant, underactuated robotic hand for dexterous grasping. In *Proc. Robotics: Science and Systems* 1687–1692 (2014).

This paper describes the design and testing of an inexpensive, modular, underactuated soft robot hand with pneumatically actuated fibre-reinforced elastomer digits.

14. Majidi, C., Shepherd, R. F., Kramer, R. K., Whitesides, G. M. & Wood, R. J. Influence of surface traction on soft robot undulation. *Int. J. Robot. Res.* **32**, 1577–1584 (2013).

15. Paul, C. Morphological computation: a basis for the analysis of morphology and control requirements. *Robot. Auton. Syst.* **54**, 619–630 (2006).

16. Hauser, H., Ijspeert, A. J., Fuchsli, R. M., Pfeifer, R. & Maass, W. Towards a theoretical foundation for morphological computation with compliant bodies. *Biol. Cybern.* **105**, 355–370 (2011).

17. Hannan, M. W. & Walker, I. D. Kinematics and the implementation of an elephant's trunk manipulator and other continuum style robots. *J. Robot. Syst.* **20**, 45–63 (2003).

18. Camarillo, D. B., Carlson, C. R. & Salisbury, J. K. Configuration tracking for continuum manipulators with coupled tendon drive. *IEEE Trans. Robot.* **25**, 798–808 (2009).

19. McMahan, W. et al. Field trials and testing of the octarm continuum manipulator. In *Proc. IEEE International Conference on Robotics and Automation* 2336–2341 (2006).

20. Lin, H., Leisk, G. & Trimmer, B. Soft robots in space: a perspective for soft robotics. *Acta Futura* **6**, 69–79 (2013).

21. Correll, N., Onal, C., Liang, H., Schoenfeld, E. & Rus, D. in *Experimental Robotics* 227–240 (Springer, 2014).

22. Jung, K. et al. Artificial annelid robot driven by soft actuators. *Bioinspir. Biomim.* **2**, S42–S49 (2007).

23. Calisti, M. et al. An octopus-bioinspired solution to movement and manipulation for soft robots. *Bioinspir. Biomim.* **6**, 036002 (2011).

24. Laschi, C. et al. Soft robot arm inspired by the octopus. *Adv. Robot.* **26**, 709–727 (2012).

25. Schulte, H. in *The Application of External Power in Prosthetics and Orthotics* 94–115 (1961).

26. Chou, C.-P. & Hannaford, B. Measurement and modeling of McKibben pneumatic artificial muscles. *IEEE Trans. Robot. Automat.* **12**, 90–102 (1996).

27. Suzumori, K., Iikura, S. & Tanaka, H. Applying a flexible microactuator to robotic mechanisms. *IEEE Control Syst.* **12**, 21–27 (1992).

28. Onal, C. D., Chen, X., Whitesides, G. M. & Rus, D. Soft mobile robots with on-board chemical pressure generation. In *Proc. International Symposium on Robotics Research* 1–16 (2011).

29. Shepherd, R. F. et al. Multigait soft robot. *Proc. Natl Acad. Sci. USA* **108**, 20400–20403 (2011).

This paper describes the rapid design and fabrication of a soft robot body capable of walking and undulating tethered to a pneumatic actuation system.

30. Marchese, A. D., Komorowski, K., Onal, C. D. & Rus, D. Design and control of a soft and continuously deformable 2D robotic manipulation system. In *Proc. IEEE International Conference on Robotics and Automation* 2189–2196 (2014).

31. Katzschmann, R. K., Marchese, A. D. & Rus, D. Hydraulic autonomous soft robotic fish for 3D swimming. In *Proc. International Symposium on Experimental Robotics* 1122374 (2014).

32. Polygerinos, P., Wang, Z., Galloway, K. C., Wood, R. J. & Walsh, C. J. Soft robotic glove for combined assistance and at-home rehabilitation. *Robot. Auton. Syst.* <http://dx.doi.org/10.1016/j.robot.2014.08.014> (2014).

This paper demonstrates a wearable soft robotic system, which augments the grasping force of wearers for rehabilitation.

33. Ilievski, F., Mazzeo, A. D., Shepherd, R. F., Chen, X. & Whitesides, G. M. Soft robotics for chemists. *Angew. Chem.* **123**, 1930–1935 (2011).

This paper demonstrates the effective use of methods and materials from chemistry and soft-materials science in the fabrication of soft robots.

34. Martinez, R. V. et al. Robotic tentacles with three-dimensional mobility based on flexible elastomers. *Adv. Mater.* **25**, 205–212 (2013).

35. Morin, S. A. et al. Camouflage and display for soft machines. *Science* **337**, 828–832 (2012).

36. Mosadegh, B. et al. Pneumatic networks for soft robotics that actuate rapidly. *Adv. Funct. Mater.* **24**, 2163–2170 (2014).

37. Park, Y.-L. et al. Design and control of a bio-inspired soft wearable robotic device for ankle-foot rehabilitation. *Bioinspir. Biomim.* **9**, 016007 (2014).

38. Bar-Cohen, Y. *Electroactive Polymer (EAP) Actuators as Artificial Muscles: Reality, Potential, and Challenges* (SPIE, 2004).

39. Wallace, G. G., Teasdale, P. R., Spinks, G. M. & Kane-Maguire, L. A. *Conductive Electroactive Polymers: Intelligent Polymer Systems* (CRC, 2008).

40. Cheng, N. G., Gopinath, A., Wang, L., Iagnemma, K. & Hosoi, A. E. Thermally tunable, self-healing composites for soft robotic applications. *Macromol. Mater. Eng.* **299**, 1279–1284 (2014).

41. Shan, W., Lu, T. & Majidi, C. Soft-matter composites with electrically tunable elastic rigidity. *Smart Mater. Struct.* **22**, 085005 (2013).

42. Steltz, E., Mozeika, A., Rodenberg, N., Brown, E. & Jaeger, H. M. JSEL: jamming skin enabled locomotion. In *Proc. International Conference on Intelligent Robots and Systems* 5672–5677 (2009).

43. Cianchetti, M. et al. Stiff-flop surgical manipulator: mechanical design and experimental characterization of the single module. In *Proc. International Conference on Intelligent Robots and Systems* 3576–3581 (2013).

44. Brown, E. et al. Universal robotic gripper based on the jamming of granular material. *Proc. Natl Acad. Sci. USA* **107**, 18809–18814 (2010).

45. Rogers, J. A., Someya, T. & Huang, Y. Materials and mechanics for stretchable electronics. *Science* **327**, 1603–1607 (2010).

46. Hammock, M. L., Chortos, A., Tee, B. C.-K., Tok, J. B.-H. & Bao, Z. The evolution of electronic skin (e-skin): a brief history, design considerations, and recent progress. *Adv. Mater.* **25**, 5997–6038 (2013).

47. Kaltenbrunner, M. et al. An ultra-lightweight design for imperceptible plastic electronics. *Nature* **499**, 458–463 (2013).

48. Yang, S. & Lu, N. Gauge factor and stretchability of silicon-on-polymer strain gauges. *Sensors* **13**, 8577–8594 (2013).

49. Kim, D.-H. et al. Epidermal electronics. *Science* **333**, 838–843 (2011).

50. Vogt, D. M., Park, Y.-L. & Wood, R. J. Design and characterization of a soft multi-axis force sensor using embedded microfluidic channels. *IEEE Sensors J.* **13**, 4056–4064 (2013).

51. Majidi, C., Kramer, R. & Wood, R. J. A non-differential elastomer curvature sensor for softer-than-skin electronics. *Smart Mater. Struct.* **20**, 105017 (2011).

52. Kramer, R. K., Majidi, C. & Wood, R. J. Masked deposition of gallium-indium alloys for liquid-embedded elastomer conductors. *Adv. Funct. Mater.* **23**, 5292–5296 (2013).

53. Muth, J. T. et al. Embedded 3D printing of strain sensors within highly stretchable elastomers. *Adv. Mater.* **26**, 6307–6312 (2014).

54. Taylor, R. F. & Schultz, J. S. *Handbook of Chemical and Biological Sensors* (CRC, 1996).

55. Wehner, M. et al. Pneumatic energy sources for autonomous and wearable soft robotics. *Soft Robotics* **1**, 263–274 (2014).

56. Goldfarb, M., Barth, E. J., Gogola, M. A. & Wehrmeyer, J. A. Design and energetic characterization of a liquid-propellant-powered actuator for self-powered robots. *IEEE/ASME Trans. Mechatronics* **8**, 254–262 (2003).

57. Shepherd, R. F. et al. Using explosions to power a soft robot. *Angew. Chem.* **125**, 2964–2968 (2013).

58. Tolley, M. T. et al. An untethered jumping soft robot. In *Proc. International Conference on Intelligent Robots and Systems* 561–566 (2014).

59. Xu, S. et al. Stretchable batteries with self-similar serpentine interconnects and integrated wireless recharging systems. *Nature Commun.* **4**, 1543 (2013).

60. Li, N., Chen, Z., Ren, W., Li, F. & Cheng, H.-M. Flexible graphene-based lithium ion batteries with ultrafast charge and discharge rates. *Proc. Natl Acad. Sci. USA* **109**, 17360–17365 (2012).

61. Suga, T., Ohshiro, H., Sugita, S., Oyaizu, K. & Nishide, H. Emerging n-type redox-active radical polymer for a totally organic polymer-based rechargeable battery. *Adv. Mater.* **21**, 1627–1630 (2009).

62. Gaikwad, A. M. et al. Highly stretchable alkaline batteries based on an embedded conductive fabric. *Adv. Mater.* **24**, 5071–5076 (2012).

63. Lipson, H. Challenges and opportunities for design, simulation, and fabrication of soft robots. *Soft Robotics* **1**, 21–27 (2014).

64. Hiller, J. & Lipson, H. Automatic design and manufacture of soft robots. *IEEE Trans. Robot.* **28**, 457–466 (2012).

65. Rieffel, J., Knox, D., Smith, S. & Trimmer, B. Growing and evolving soft robots. *Artif. Life* **20**, 143–162 (2014).

66. Cho, K. J. et al. Review of manufacturing processes for soft biomimetic robots. *Int. J. Precis. Eng. Man.* **10**, 171–181 (2009).

67. Lipson, H. & Kurman, M. *Fabricated: The New World of 3D Printing* (Wiley, 2013).

68. Cham, J. G., Bailey, S. A., Clark, J. E., Full, R. J. & Cutkosky, M. R. Fast and robust hexapedal robots via shape deposition manufacturing. *Int. J. Robot. Res.* **21**, 869–882 (2002).

69. Xia, Y. & Whitesides, G. M. Soft lithography. *Annu. Rev. Mater. Sci.* **28**, 153–184 (1998).

70. Marchese, A. D., Katzschmann, R. & Rus, D. A recipe for soft fluidic elastomer robots. *Soft Robotics* **2**, 7–25 (2015).

71. Sumbre, G., Fiorito, G., Flash, T. & Hochner, B. Octopus uses human-like strategy to control point-to-point arm movement. *Curr. Biol.* **16**, 767–772 (2006).

72. Margheri, L., Laschi, C. & Mazzolai, B. Soft robotic arm inspired by the octopus: I. from biological functions to artificial requirements. *Bioinspir. Biomim.* **7**, 025004 (2012).

73. Lin, H.-T., Leisk, G. G. & Trimmer, B. GoQBot: a caterpillar-inspired soft-bodied rolling robot. *Bioinspir. Biomim.* **6**, 026007 (2011).

This paper describes a soft mobile robot developed as a tool to study caterpillar locomotion including crawling and ballistic rolling.

74. Saunders, F., Trimmer, B. A. & Rife, J. Modeling locomotion of a soft-bodied arthropod using inverse dynamics. *Bioinspir. Biomim.* **6**, 016001 (2011).

75. Webster, R. J. & Jones, B. A. Design and kinematic modeling of constant curvature continuum robots: a review. *Int. J. Robot. Res.* **29**, 1661–1683 (2010).

76. Gravagne, I. A., Rahn, C. D. & Walker, I. D. Large deflection dynamics and control for planar continuum robots. *IEEE/ASME Trans. Mechatronics* **8**, 299–307 (2003).

77. Jones, B. A. & Walker, I. D. Kinematics for multisegment continuum robots. *IEEE Trans. Robot.* **22**, 43–55 (2006).

78. Renda, F., Giorelli, M., Calisti, M., Cianchetti, M. & Laschi, C. Dynamic model of a multi-bending soft robot arm driven by cables. *IEEE Trans. Robot.* **30**, 1109–1122 (2014).

79. Neppalli, S., Csencsits, M. A., Jones, B. A. & Walker, I. D. Closed-form inverse

kinematics for continuum manipulators. *Adv. Robot.* **23**, 2077–2091 (2009).

80. Wang, H. et al. Visual servo control of cable-driven soft robotic manipulator. In *Proc. International Conference on Intelligent Robots and Systems* 57–62 (2013).
81. Khatib, O., Sentis, L., Park, J. & Warren, J. Whole body dynamic behavior and control of human-like robots. *Int. J. Humanoid Robot.* **1**, 29–43 (2004).
82. Napp, N., Araki, B., Tolley, M. T., Nagpal, R. & Wood, R. J. Simple passive valves for addressable pneumatic actuation. In *Proc. International Conference on Robotics and Automation* 1440–1445 (2014).
83. Chirikjian, G. S. Hyper-redundant manipulator dynamics: a continuum approximation. *Adv. Robot.* **9**, 217–243 (1994).
84. Yekutieli, Y. et al. Dynamic model of the octopus arm. I. biomechanics of the octopus reaching movement. *J. Neurophysiol.* **94**, 1443–1458 (2005). **This paper describes a 2D dynamic model for a soft manipulator based on muscular hydrostats.**
85. Snyder, J. & Wilson, J. Dynamics of the elastic with end mass and follower loading. *J. Appl. Mech.* **57**, 203–208 (1990).
86. Tatlicioğlu, E., Walker, I. D. & Dawson, D. M. Dynamic modeling for planar extensible continuum robot manipulators. In *Proc. International Conference on Robotics and Automation* 1357–1362 (2007).
87. Luo, M., Agheli, M. & Onal, C. D. Theoretical modeling and experimental analysis of a pressure-operated soft robotic snake. *Soft Robotics* **1**, 136–146 (2014).
88. Seok, S. et al. Meshworm: a peristaltic soft robot with antagonistic nickel titanium coil actuators. *IEEE/ASME Trans. Mechatronics* **18**, 1485–1497 (2013).
89. Stokes, A. A., Shepherd, R. F., Morin, S. A., Ilievski, F. & Whitesides, G. M. A hybrid combining hard and soft robots. *Soft Robotics* **1**, 70–74 (2014).
90. Amend, J. R., Brown, E. M., Rodenberg, N., Jaeger, H. M. & Lipson, H. A positive pressure universal gripper based on the jamming of granular material. *IEEE Trans. Robot.* **28**, 341–350 (2012).
91. Sanan, S., Lynn, P. S. & Griffith, S. T. Pneumatic torsional actuators for inflatable robots. *J. Mech. Robot.* **6**, 031003 (2014).
92. Kramer, R. K., Majidi, C. & Wood, R. J. Wearable tactile keypad with stretchable artificial skin. In *Proc. International Conference on Robotics and Automation* 1103–1107 (2011).
93. Mengüç, Y. et al. Wearable soft sensing suit for human gait measurement. *Inter. J. Robotics Res.* **33**, 1748–1764 (2014).
94. Song, Y. S. et al. Soft robot for gait rehabilitation of spinalized rodents. In *Proc. International Conference on Intelligent Robots and Systems* 971–976 (2013).
95. Roche, E. T. et al. A bioinspired soft actuated material. *Adv. Mater.* **26**, 1200–1206 (2014).
96. Ieropoulos, I., Anderson, I. A., Gisby, T., Wang, C.-H. & Rossiter, J. Microbial-powered artificial muscles for autonomous robots. *Proc. SPIE* **7287**, 728708–728708 (2009).
97. Nawroth, J. C. et al. A tissue-engineered jellyfish with biomimetic propulsion. *Nature Biotechnol.* **30**, 792–797 (2012).
98. Chambers, L., Winfield, J., Ieropoulos, I. & Rossiter, J. Biodegradable and edible gelatine actuators for use as artificial muscles. *Proc. SPIE* **9056**, 90560B (2014).
99. Pfeifer, R., Bongard, J. & Grand, S. *How the Body Shapes The Way We Think: A New View of Intelligence* (MIT Press, 2007).
100. Suzumori, K., Endo, S., Kanda, T., Kato, N. & Suzuki, H. A bending pneumatic rubber actuator realizing soft-bodied manta swimming robot. In *Proc. International Conference on Robotics and Automation* 4975–4980 (2007).

Acknowledgements The authors would like to thank A. Marchese and R. Katzschmann for their insightful comments and suggestions in editing this paper. This work was done with partial support from the National Science Foundation grant number IIS-1226883, for which we are grateful.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of this paper at go.nature.com/dssyt. Correspondence should be addressed to D.R. (rus@csail.mit.edu).

From evolutionary computation to the evolution of things

Agoston E. Eiben¹ & Jim Smith²

Evolution has provided a source of inspiration for algorithm designers since the birth of computers. The resulting field, evolutionary computation, has been successful in solving engineering tasks ranging in outlook from the molecular to the astronomical. Today, the field is entering a new phase as evolutionary algorithms that take place in hardware are developed, opening up new avenues towards autonomous machines that can adapt to their environment. We discuss how evolutionary computation compares with natural evolution and what its benefits are relative to other computing approaches, and we introduce the emerging area of artificial evolution in physical systems.

The proposal that evolution could be used as a metaphor for problem solving came with the invention of the computer¹. In the 1970s and 1980s the principal idea was developed into different algorithmic implementations under names such as evolutionary programming², evolution strategies^{3,4} and genetic algorithms⁵, followed later by genetic programming⁶. These branches merged in the 1990s, and in the past 20 years so-called evolutionary computation or evolutionary computing has proven to be highly successful across a wide range of computational tasks in optimization, design and modelling^{7–9}. For instance, urgent needs in the development of low-cost thin-film photovoltaic technologies were addressed using genetic algorithms for topology optimization¹⁰. This led to highly efficient light-trapping structures that exhibited more than a threefold increase over a classic limit, and achieved efficiency levels far beyond the reach of intuitive designs. It has also been convincingly demonstrated that evolutionary approaches are powerful methods for knowledge discovery. For example, equations were evolved to model motion-tracking data captured from various physical systems, ranging from simple harmonic oscillators to chaotic double pendula. This approach discovered, with limited prior knowledge of physics, kinematics or geometry, several laws of geometric and momentum conservation, and uncovered the ‘alphabet’ used to describe those systems¹¹.

From the perspective of the underlying substrate in which the evolution takes place, the emergence of evolutionary computation can be considered as a major transition of the evolutionary principles from ‘wetware’, the realm of biology, to software, the realm of computers. Today the field is at an exciting stage. New developments in robotics and rapid prototyping (3D printing) are paving the way towards a second major transition: from software to hardware, going from digital evolutionary systems to physical ones^{12,13}.

In this Review we outline the working principles of evolutionary algorithms, and briefly discuss the differences between artificial and natural evolution. We illustrate the power of evolutionary problem solving by discussing a number of successful applications, reflect on the features that make evolutionary algorithms so successful, review the current trends of the field, and give our perspective on future developments.

Evolution and problem solving

The essence of an evolutionary approach to solve a problem is to equate possible solutions to individuals in a population, and to introduce a notion of fitness on the basis of solution quality. To obtain a working

evolutionary algorithm one has to go through a number of design steps. The first step is to identify a representation: a suitable data structure that can represent possible solutions to the problem. The next step is to define a way of measuring the quality of an individual based on problem-specific requirements. The final step is to specify suitable selection and variation operators (Fig. 1).

Analogous to natural evolution, an evolutionary algorithm can be thought of as working on two levels. At the higher level (the original problem context), phenotypes (candidate solutions) have their fitness measured. Selection mechanisms then use this measure to choose a pool of parents for each generation, and decide which parents and offspring go forward to the next generation. At the lower level, genotypes are objects that represent phenotypes in a form that can be manipulated to produce variations (Box 1). Genotype–phenotype mapping bridges the two levels. At the genotypic level, variation operators generate new individuals (offspring) from selected parents. Mutation operators are based on one parent (asexual reproduction) and randomly change some values. Recombination operators create offspring by combining values from the genotypes of two (or more) parents. Finally, an execution manager controls the overall functioning of the algorithm. It regulates the initialization of the first population, the execution of the selection–variation cycles, and the termination of the algorithm. It also manages the population size (typically kept constant) and other parameters affecting selection and variation. For instance, it determines the number of parents per generation, and whether mutation, recombination or both produce the offspring for a given set of parents.

Evolutionary algorithms are easily transferable from one application to another because only two components are problem dependent: the way that the genotypes are converted to phenotypes and the fitness function. The history of evolutionary computation has shown that suitable combinations of a few simple data structures can represent possible solutions to a huge variety of different problems (Box 1). In other words, a relatively small collection of possible genotypes can accommodate many different kinds of phenotypes. Just as the genetic mechanisms underpinning natural evolution are largely species independent, acting on DNA or RNA, so too in evolutionary computation the choice of suitable variation operators depends solely on the data structure present in the genotypes and not on the specific problem being tackled. Selection operators do not even depend on the chosen representation, as they only consider fitness information. This implies that for a certain problem a suitable evolutionary algorithm can be designed easily, as long as the

¹VU University Amsterdam, de Boelelaan 1081a, 1081HV Amsterdam, the Netherlands. ²University of the West of England, Bristol BS16 1QY, UK.

problem-dependent phenotypes can be mapped to one of the ‘standard’ genotypes. From that point on, freely available evolutionary algorithm machinery can be used.

It should be noted that just because an algorithm is formally suitable, it does not necessarily mean it will be successful. Suitability only means that the evolutionary algorithm is capable of searching through the space of possible solutions of a problem, but gives no guarantees that this search will be either effective or efficient.

Positioning of evolutionary computation

From a historical perspective, humans have had two roles in evolution. Just like any other species, humans are the product of, and are subject to, evolution. But for millennia (in fact, for about twice as long as we have used wheels) people have also actively influenced the course of evolution in other species — by choosing which plants or animals should survive or mate. Thus humans have successfully exploited evolution to create improved food sources¹⁴ or more useful animals¹⁵, even though the mechanisms involved in the transmission of traits from one generation to the next were not understood.

Historically, the scope of human influence in evolution was very limited, being restricted to interfering with selection for survival and reproduction. Influencing other components, such as the design of genotypes, or mutation and recombination mechanisms, was far beyond our reach. This changed with the invention of the computer, which provided the possibility of creating digital worlds that are very flexible and much more controllable than the physical reality we live in. Together with the increased understanding of the genetic mechanisms behind evolution, this brought about the opportunity to become active masters of evolutionary processes that are fully designed and executed by human experimenters ‘from above’.

It could be argued that evolutionary algorithms are not faithful models of natural evolution (Table 1). However, they certainly are a form of evolution. As Dennett¹⁶ said “If you have variation, heredity, and selection, then you must get evolution”.

From a computer-science perspective, evolutionary algorithms are

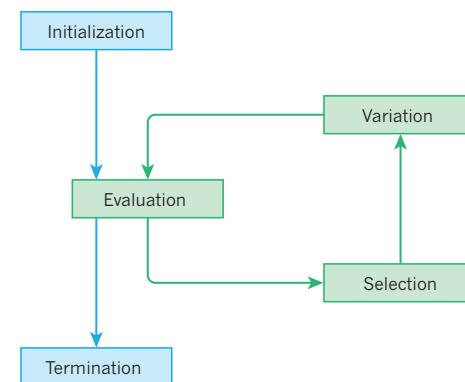


Figure 1 | The principal diagram of evolutionary algorithms. The initialization process seeds the search with a population of randomly created solutions. After this the algorithm enters a loop of evaluating the current generation of solutions, selecting some to act as the basis for the next generation, and then creating new solutions through variation (mutation or crossover). Periodically, the algorithm checks to see whether user-specified termination criteria are met — such as reaching a desired level of fitness, or undergoing a certain number of generations without improvement.

randomized heuristic search methods based on the generate-and-test principle: producing an offspring amounts to generating a new point in the search space, and testing is done through fitness evaluation. What distinguishes evolutionary algorithms from other algorithms in computer science is the unique combination of being stochastic and maintaining their working memory in the form of a population of candidate solutions. It should be noted that there are many variations of the generic evolutionary computation template under various names. Today, the family of evolutionary algorithms includes historical members: genetic algorithms, evolution strategies, evolutionary programming and genetic programming; and younger siblings, such as differential evolution and particle swarm optimization^{17–25}. These

BOX 1

Representing solutions to problems as genotypes

Genotypes are objects that represent phenotypes (candidate solutions to problems) in a form that can be manipulated to produce variations.

Examples of data structures frequently used as genotypes are shown in the Figure. One suitable mutation operator is shown for each, with its action shown by red arrows. Note that the mutation operator must deliver a child of the same data type — for example, a valid mutation operator for permutations must result in a valid permutation. Complex problems might require complex genotypes with appropriate mutation operators.

Bitstrings are the natural choice for problems for which solutions are composed of on/off or true/false decisions. The most commonly used mutation operator makes an independent choice in each position whether to invert the bit value.

Permutations can be used when the problem involves ordering a set of entities, such as in routing or scheduling. One simple mutation operator swaps the values in two randomly selected locations.

Real-valued vectors can capture continuous optimization problems, for example where the variables represent quantities such as dimensions or mass. Typically, the mutation operator perturbs each value by adding a (small) random number.

Trees are branching data structures suitable for representing equations, logical expressions or program code. A common mutation operator selects a node at random, and replaces the subtree below with a new, randomly generated one.

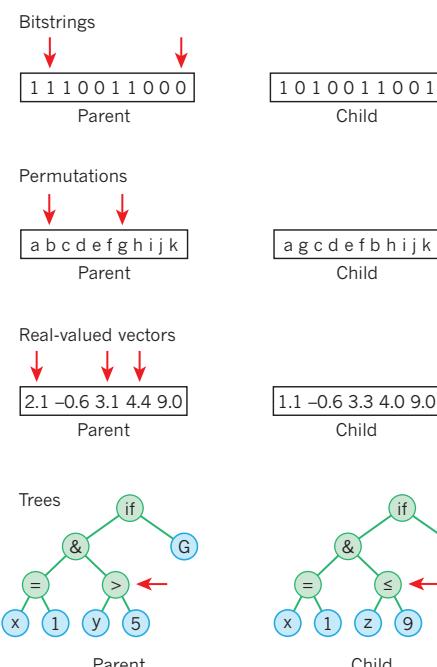


Table 1 | Main differences between natural evolution and evolutionary algorithms

Natural evolution	Evolutionary algorithms
Fitness Observed quantity: <i>a posteriori</i> effect of selection and reproduction ('in the eye of the observer').	Predefined <i>a priori</i> quantity that drives selection and reproduction.
Selection Complex multifactor force based on environmental conditions, other individuals of the same species and those of other species (predators). Viability is tested continually; reproducibility is tested at discrete times.	Randomized operator with selection probabilities based on given fitness values. Survivor selection and parent selection both happen at discrete times.
Genotype–phenotype mapping Highly complex biochemical and developmental process influenced by the environment.	Typically a simple mathematical transformation or parameterized procedure. A few systems use generative and developmental genotype–phenotype maps.
Variation Offspring are created from one (asexual reproduction) or two parents (sexual reproduction). Horizontal gene transfer can accumulate genes from more individuals.	Unconstrained vertical gene transfer. Offspring may be generated from any number of parents: one, two or many.
Execution Parallel, decentralized execution; birth and death events are not synchronized.	Typically centralized with synchronized birth and death.
Population Spatial embedding implies structured populations. Population size varies according to the relative number of birth and death events. Populations can and do go extinct.	Typically unstructured and panmictic (all individuals are potential partners). Population size is usually kept constant by synchronizing time and number of birth and death events.

differ in some details, terminology or motivational metaphor, but are, in essence, all instances of the same algorithmic template.

It is common to categorize algorithms according to completeness (can they generate every possible solution), optimality (are they guaranteed to find the best solution, and identify it as such) and efficiency. The completeness of an evolutionary algorithm can be achieved by an appropriate choice of representation and variation operators. The optimality is a more complex issue. Although optimal methods exist for many problems, their run time scales so poorly that they are impractical to use in most non-trivial cases — hence the interest in heuristic methods. As long as the heredity principle (similar individuals have similar fitness) holds, an evolutionary algorithm will have a ‘basic instinct’ to improve the population’s fitness over time — because the selection operators are biased towards choosing fitter individuals for reproduction and survival. Thus, if we can define artificial fitness on the basis of a criterion grounded in the problem to be solved then the evolutionary algorithm will tend to find solutions that optimize the fitness values, or at least approximate them. This implies that evolutionary algorithms can be used to solve optimization problems and, consequently, any problem that can be transformed into an equivalent optimization task. This includes most problems in design, and those connected with building or learning models from data. Nevertheless, it is important to understand that evolutionary algorithms are not optimizers²⁶, but approximators, and they are not optimal since we might not know whether the fitness of the best evolved solution is in fact the highest value possible. Yet, they become very interesting when approximate solutions are acceptable, for instance, if the global optimum is not known or not required.

Applications of evolutionary computation

The hypothesis that embedding the principles of evolution within computer algorithms can create powerful mechanisms for solving difficult, poorly understood problems is now supported by a huge body of evidence. Evolutionary problem solvers have proven capable of delivering high-quality solutions to difficult problems in a variety of scientific and technical domains, offering several advantages over conventional optimization and design methods.

One appealing example from the design domain concerns X-band antennas for the NASA Space Technology 5 (ST5) spacecraft²⁷. The normal approach to this task is very time and labour intensive, relying heavily on expert knowledge. The evolutionary-algorithm-based

approach not only discovered effective antenna designs, but could also adjust designs quickly when requirements changed. One of these antennas was actually constructed and deployed on the ST5 spacecraft, thus becoming the first computer-evolved hardware in space. This project also demonstrates a specific advantage of evolutionary over manual design. The evolutionary algorithms generated and tested thousands of completely new solutions, many with unusual structures that expert antenna designers would be unlikely to produce. Evolutionary algorithms have also been successful in many other aeronautical and aerospace engineering endeavours. Problems in this field typically have highly complex search spaces and multiple conflicting objectives. Population-based methods such as evolutionary algorithms have proven effective at meeting the challenges of this combination. In particular, so-called multi-objective evolutionary algorithms change the selection function to explicitly reward diversity, so that they discover and maintain high-quality solutions representing different trade-offs between objectives — technically, they approximate diverse segments of the Pareto front²⁸. Many examples can also be found in bioinformatics. For instance, by mining the ChEMBL database (which contains bioactive molecules with drug-like properties), a set of transformations of chemical structures was identified that were then used as the mutation operator in an automated drug-design application²⁹. The results showed clear benefits, particularly in accommodating multiple target profiles such as desired polypharmacology. This nicely illustrates how other approaches, or existing knowledge, can be easily co-opted or accommodated within an evolutionary computing framework.

Numerical and combinatorial optimization are important application areas of evolutionary algorithms. Particularly challenging is black-box optimization, where the nature of the objective function requires numerical (rather than analytical) methods, and gradient information can only be approximated by sampling solutions. A systematic experimental study compared mathematical programming and evolutionary computation of a range of synthetic black-box optimization problems, which were allowed different amounts of computing time and resources³⁰. The results showed that mathematical programming algorithms — that were designed to provide quick progress in the initial stages of the search — outperform evolutionary algorithms if the maximum number of evaluations is low, but this picture changes if the computational budget is increased. Ultimately, the study concludes that an evolutionary algorithm, in particular BIPOP-CMA-ES, is able to find the optimum of a broader class of functions, solve problems with

a higher precision and solve some problems faster. The power of evolution strategies (especially the very successful CMA-ES variants³¹) for real-life black-box optimization problems from industry has been discussed extensively³². Evidence gathered from years of academic research and development for industrial applications suggests that the niche for evolution strategies is formed by optimization tasks with a very limited budget for how many solutions can have their fitness evaluated. Although this finding is not in line with conventional wisdom within the field, there is ample support for this proposal.

Machine learning and modelling is another prominent area in which evolutionary algorithms have proved their power, especially as many contemporary approaches would otherwise rely on (often crude) greedy or local search algorithms to refine and optimize models. For example, neuroevolutionary approaches use evolutionary algorithms to optimize the structure, parameters, or both simultaneously, of artificial neural networks^{33,34}. In other branches of machine learning, using evolutionary computing to design algorithms has been shown to be very effective as an alternative to handcrafting them, for instance, for inducing decision trees³⁵. Furthermore, evolutionary algorithms have been applied to prediction problems. For instance, to tackle the problem of predicting the tertiary structure of a protein, an algorithm was designed to evolve a key component of automated predictors — the function used to estimate a structure's energy³⁶. State-of-the-art methods in protein-structure prediction are limited by assuming a linear combination of energy terms, whereas a genetic programming method easily accommodates expressions based on a much richer syntax. The best energy function found by the genetic programming algorithm provided significantly better prediction guidance than conventional functions. The algorithm was able to automatically discover the most and least useful energy terms, without having any knowledge of how these terms alone are correlated to the prediction error.

The design of controllers for physical entities, such as machinery or robots, has proved to be another fruitful area. For example, control strategies for operating container cranes were evolved using a physical crane to determine fitness values³⁷. The evolution of controllers is also possible *in situ*, for example in a population of robots during, and not just before, their operational period^{38,39}. Evolutionary robotics⁴⁰ is an especially challenging application area because of two additional issues that other branches of evolutionary computing do not face: the very weak and noisy link between controllable design details and the target feature or features; and the great variety of conditions under which a solution should perform well. Normally in evolutionary computing there is a three-step evaluation chain: genotype to phenotype to fitness. For robots the chain is four-step: genotype to phenotype to behaviour to fitness. In this four-step chain the robots morphology and controller form the phenotype. However, it could be argued that the behaviour should be considered as phenotype, because it is the entity that is being evaluated. Furthermore, the behaviour depends on many external factors, creating an unpredictable environment in which the robot is expected to perform. Nevertheless, since the manual design of an autonomous and adaptive mobile robot is extremely difficult, evolutionary approaches offer large potential benefits. These include the possibility of continuous and automated design, manufacture and deployment of robots of very different morphologies and control systems⁴¹. Several studies have demonstrated such benefits, in which robot control systems that were automatically generated by artificial evolution were comparatively simpler or more efficient than those engineered using other design methods⁴². In all cases, robots initially exhibited uncoordinated behaviour, but a few hundreds of generations were sufficient to achieve efficient behaviours in a wide range of experimental conditions.

Several state-of-the-art algorithms for applications across a great variety of problem domains are based on hybridizing evolutionary search with existing algorithms, especially local search methods. This kind of hybridization can be thought of as adding 'lifetime learning' to the evolutionary process. Freed from the restrictions of natural evolution (such as

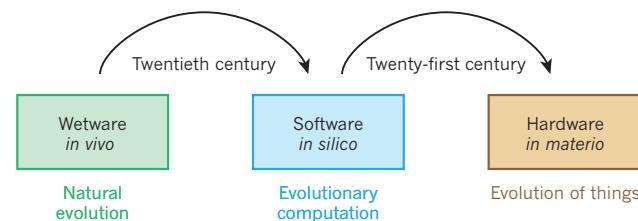


Figure 2 | Two major transitions in the history of artificial evolution. In the twentieth century computer technology enabled artificial Darwinian processes *in silico* — the evolution of digital entities. In the twenty-first century, developments in robotics, materials science and 3D printing will enable the evolution of physical artefacts or machines.

learned traits not being written back immediately to the genotype), and being able to experiment with novel types of individual and social learning, the theory and practice of so-called memetic algorithms has become an important topic in the field^{43–45}. Such hybrid algorithms can often find good (or better) solutions faster than a pure evolutionary algorithm when the additional method searches systematically in the vicinity of good solutions, rather than relying on the more randomized search carried out by mutation^{46,47}. For example, the cell suppression problem (deciding which data cells to disclose in published statistical tables in order to protect respondents' confidentiality)⁴⁸ was solved using a combination of graph partitioning, linear programming and evolutionary optimization of the sequence in which vulnerable cells were considered. This produced methods that could protect published statistical tables at a size that was several orders of magnitude greater than had previously been possible. Memetic algorithms have obtained an eminent place among the best approaches to solving really hard problems.

State of the art

Although initially considerable scepticism surrounded evolutionary algorithms, over the past 20 years evolutionary computation has grown to become a major field in computational intelligence^{7–9}. As well as solving hard problems in various application areas, the emphasis of evolutionary algorithms on randomness as a source of variation has been shown to have particular advantages: the lack of problem-specific preconceptions and biases of the algorithm designer opens up the way to unexpected 'original' solutions that can even have artistic value^{49–51} (<http://endlessforms.com/>). The perception of evolution as a problem solver has broadened from seeing evolution as a heuristic algorithm for (parametric) optimization to considering it to be a powerful approach for (structural) design^{52,53}.

In general, evolutionary algorithms have proven competitive in solving hard problems in the face of challenging characteristics like non-differentiability, discontinuities, multiple local optima, noise and nonlinear interactions among the variables, especially if the computational budgets are sufficiently high. Evolution is a slow learner, but the steady increase in computing power, and the fact that the algorithm is inherently suited to parallelization, mean that more and more generations can be executed within practically acceptable timescales.

The performance of evolutionary algorithms has also been compared with that of human experts, and there is now substantial and well-documented evidence of evolutionary algorithms producing measurably human-competitive results⁵⁴. The annual Humies competition (<http://www.genetic-programming.org/combined.php>), which rewards human-competitive results from evolutionary computation, highlights the great variety of hard problems for which evolutionary algorithms have delivered excellent solutions.

The success and popularity of evolutionary algorithms can be attributed to a number of algorithmic features, which makes them attractive. First, they are assumption free because applying an evolutionary algorithm consists of specifying the representation for candidate solutions and providing an external function that first transforms the genotype into a candidate solution and then provides an evaluation. Internally,

evolutionary algorithms make no explicit assumptions about the problem, hence they are widely applicable and easily transferable at low cost. Second, they are flexible because they can be easily used in collaboration with existing methods, such as local search; they can be incorporated within, or make use of, existing toolsets; and combinations with domain-specific methods often lead to superior solvers because they can exploit the best features of different approaches. Third, they are robust, owing to the use of a population and randomized choices, which mean evolutionary algorithms are less likely to get trapped in suboptimal solutions than other search methods. They are also less sensitive to noise or infidelity in the models of the system used to evaluate solutions, and can cope with changes in the problem. Fourth, they are not focussed on a single solution because having a population means that an algorithm terminates with a number of solutions. Thus, users do not have to pre-specify their preferences and weighting in advance, but can make decisions after they see what is possible to achieve. This is a great advantage for problems with many local optima, or with a number of conflicting objectives. Finally, they are capable of producing unexpected solutions because they are blind to human preconceptions and so can find effective, but non-intuitive solutions, which are often valuable in design domains.

The theoretical underpinning of evolutionary algorithms remains a hard nut to crack. Mathematical analysis can illuminate some properties, but even digital evolutionary processes exhibit very complex dynamics that allow only limited theory forming, despite the diverse set of tools and methods ranging from quantitative genetics to statistical physics⁵⁵. One important theoretical result is the no free lunch theorem. This states that evolutionary algorithms are not generic super solvers — but neither is any other method, because there is no such thing⁵⁶. Instead, “an evolutionary algorithm is the second best solver for any problem”, meaning that in many cases a carefully hand-crafted solver that exploits problem characteristics is superior for the problem at hand, but that it might take years to create that solver. A long-standing issue for theorists is algorithm convergence. Early results were based on Markov-chain analysis and addressed convergence in general⁵⁷, but more recent work found specific relationships between algorithmic set-up and expected run times⁵⁸. Despite all the difficulties, the field is making progress in theory^{59,60}.

Important research trends

The development of evolutionary computation continues along a number of research threads.

Automated design and tuning of evolutionary algorithms

Experience has shown that there are several design choices behind an evolutionary algorithm that greatly influence its performance. To reduce the number of design decisions to be made, and the impact of poor choices, the community is working on automated design aids. These can customize an initial algorithm set-up for a given problem offline (before the run) or online (during the run)⁶¹. Techniques such as automated parameter tuning^{62–65} and adaptive parameter control continue to make advances in this area^{66–69}.

Using surrogate models

Increasingly, evolutionary algorithms are being used for problems in which evaluating each population member over many generations would take too long to permit effective evolution given the resources available. A range of approaches — collectively known as surrogate models — are being developed that use computationally cheaper models in place of full fitness evaluations, and that refine those models through occasional full evaluations of targeted individuals^{70–73}.

Handling many objectives

Having proven highly successful for finding solutions to problems with multiple objectives (typically up to ten)⁷⁴, the community is now making rapid advances in the field of many objectives — moving way beyond the capabilities of other algorithms^{75–77}. In tandem with

algorithmic advances, this has spurred renewed interest in interactive evolutionary algorithms, which have been successfully applied to elicit user preferences and knowledge in many areas from design to art⁵¹ (<http://picbreeder.org>). Results suggest a useful synergy with periodic user interaction to incorporate preferences that help to focus the search down to a more manageable set of dimensions⁷⁸. Importantly, this involves eliciting user preferences in response to what is discovered to be possible, rather than *a priori*.

Generative and developmental representations

Further to the conventionally simple genotype–phenotype mappings, the use of indirect encodings is gaining traction. Such generative and developmental representations allow the reuse of code, which helps to scale up the complexity of artificially evolved phenotypes, for instance, in evolutionary robotics, artificial life and morphogenetic engineering^{79–83}.

Outlook

The range of problems to which evolutionary algorithms have been successfully applied has grown year on year, and there is every reason to expect this to continue. In the future, we expect to see increasing interest in applying evolutionary algorithms to embodied or embedded systems; that is, employing evolution in populations for which the candidate solutions are controllers or drivers that implement the operational strategy for some situated entities, and are evaluated within the context of some rich dynamic environment; not for what they are, but for what they do. Examples include policies for Web-crawlers, information retrieval strategies, software for machinery and smart devices, and controllers for autonomous robots^{84,85}. In such cases the evolved solutions are embedded in entities that exist and act in a ‘habitat’, the internet or the physical world, that is too complex and dynamic to be modelled perfectly. Enhancing the system with the ability to evolve and adapt after deployment can complement the offline optimization approach employed during the design stage. The novelty of such systems is that evolutionary changes take place within the operational period. These systems will be different because they replace the conventional design-and-deploy approach by a design–deploy–adapt loop in which the evolutionary component is a principal part of the system.

This approach is already gaining traction in two areas. In the field of search-based software engineering, evolutionary algorithms are gaining prominence in response to the mismatch between the availability of expert software engineers and the explosion of interconnected devices requiring new and/or updated software⁸⁶. Meanwhile, recent developments in rapid fabrication technologies (3D printing) and ever smaller and more powerful robotic platforms mean that evolutionary computing is now starting to make the next major transition to the automated creation of physical artefacts and ‘smart’ objects⁸⁷ (Fig. 2). In the long term this could lead to a disruptive robotic technology in which design and production are replaced by selection and reproduction without the involvement of human designers and human-operated facilities.

Last but not least, we foresee a fruitful cross-fertilization with biology in the coming decade based on a bidirectional flow of inspiration, understanding and knowledge. On the one hand, the advancing insights in molecular and evolutionary biology can be used to make more sophisticated evolutionary algorithms and may help to solve previously intractable problems. The opportunities and challenges of this avenue have been outlined in a research agenda to transform artificial evolution to computational evolution⁸⁸. On the other hand, a new kind of artificial evolution — the evolution of things — opens new horizons for biologists. In 1992, the evolutionary biologist John Maynard Smith commented: “So far, we have been able to study only one evolving system and we cannot wait for interstellar flight to provide us with a second. If we want to discover generalizations about evolving systems, we will have to look at artificial ones”⁸⁹. Artificial evolution implemented on real hardware, as in evolutionary robotics, offers a new research instrument to this end^{42,90–95}. The use of real hardware overcomes the principal

deficiency of software models, which lack the richness of matter that is a source of challenges and opportunities not yet matched in artificial algorithms⁹⁶. Hence, they can provide new insights into fundamental issues such as the factors influencing evolvability, resilience, the rate of progress under various circumstances, or the co-evolution of mind and body. Using a non-biochemical substrate for such research is becoming technologically ever more feasible, and it increases the generalizability of the findings. In particular, using a different medium for evolutionary studies can separate generic principles and ground truth from effects that are specific for carbon-based life as we know it. ■

Received 18 December 2014; accepted 24 March 2015.

- Turing, A. M. in *Machine Intelligence 5* (eds Meltzer, B. & Michie, D.) (Edinburgh Univ. Press, 1969).
- Fogel, L. Owens, A. J. & Walsh, M. J. *Artificial Intelligence Through Simulated Evolution* (Wiley, 1966).
- Rechenberg, I. *Evolutionstrategie: Optimierung Technischer Systeme nach Prinzipien des Biologischen Evolution* [in German] (Fromman-Hozlboog, 1973).
- Schwefel, H.-P. *Numerical Optimization of Computer Models* (Birkhäuser, 1977).
- Holland, J. H. *Adaption in Natural and Artificial Systems* (Univ. Michigan Press, 1975).
- Koza, J. R. *Genetic Programming* (MIT Press, 1992).
- Eiben, A. E. & Smith, J. E. *Introduction to Evolutionary Computing* (Springer, 2003).
- Ashlock, D. *Evolutionary Computation for Modeling and Optimization* (Springer, 2006).
- De Jong, K. *Evolutionary Computation: a Unified Approach* (MIT Press, 2006).
- Wang, C., Yu, S., Chen, W. & Sun, C. Highly efficient light-trapping structure design inspired by natural evolution. *Sci. Rep.* **3**, 1025 (2013).
- Schmidt, M. & Lipson, H. Distilling free-form natural laws from experimental data. *Science* **324**, 81–85 (2009).
- This paper provides a forceful demonstration of the power of evolutionary methods for tasks that are thought to require highly educated scientists to perform.**
- Eiben, A. E., Kernbach, S. & Haasdijk, E. Embodied artificial evolution: artificial evolutionary systems in the 21st Century. *Evol. Intel.* **5**, 261–272 (2012).
- Eiben, A. E. in *Parallel Problem Solving from Nature – PPSN XII* (eds Filipic, B., Bartz-Beielstein, T., Branke, J. & Smith, J.), 24–39 (Springer, 2014).
- Piperno, D. R., Ranere, A. J., Holst, I., Irarite, J. & Dickau, R. Starch grain and phytolith evidence for early ninth millennium B.P. maize from the Central Balsas River Valley, Mexico. *Proc. Natl Acad. Sci. USA* **106**, 5019–5024 (2009).
- Akey, J. M. et al. Tracking footprints of artificial selection in the dog genome. *Proc. Natl Acad. Sci. USA* **107**, 1160–1165 (2010).
- Dennett, D. *Darwin's Dangerous Idea* (Penguin, 1995).
- Goldberg, D. *Genetic Algorithms in Search, Optimization, and Machine Learning* (Addison-Wesley, 1989).
- Fogel, D.B. *Evolutionary Computation* (IEEE, 1995).
- Schwefel, H.-P. *Evolution and Optimum Seeking* (Wiley, 1995).
- Bäck, T. *Evolutionary Algorithms in Theory and Practice* (Oxford Univ. Press, 1996).
- Banzhaf, W., Nordin, P., Keller, R. E. & Francone, F. D. *Genetic Programming: an Introduction* (Morgan Kaufmann, 1998).
- Storn, R. & Price, K. Differential evolution — a simple and efficient heuristic for global optimization over continuous spaces. *J. Glob. Optim.* **11**, 341–359 (1997).
- Price, K. V., Storn, R. N. & Lampinen, J. A. *Differential Evolution: a Practical Approach to Global Optimization* (Springer, 2005).
- Kennedy, J. & Eberhart, R. C. Particle swarm optimization. In *Proc. IEEE International Conference on Neural Networks 1942–1948* (IEEE, 1995).
- Kennedy, J. & Eberhart, R.C. *Swarm Intelligence* (Morgan Kaufmann, 2001).
- De Jong, K. A. Are genetic algorithms function optimizers? In *Proc. 2nd Conference on Parallel Problem Solving from Nature* (eds Manner, R. & Manderick, B.), 3–13 (North-Holland, 1992).
- Hornby, G. S., Lohn, J. D. & Linden, D. S. Computer-automated evolution of an X-band antenna for NASA's space technology 5 mission. *Evol. Comput.* **19**, 1–23 (2011).
- Arias-Montano, A., Coello, C. A. C. & Mezura-Montes, E. Multiobjective evolutionary algorithms in aeronautical and aerospace engineering. *IEEE Trans. Evol. Comput.* **16**, 662–694 (2012).
- Besnard, J. et al. Automated design of ligands to polypharmacological profiles. *Nature* **492**, 215–220 (2012).
- Posik, P., Huyer, W. & Pal, L. A comparison of global search algorithms for continuous black box optimization. *Evol. Comput.* **20**, 509–541 (2012).
- Hansen, N. & Ostermeier, A. Completely derandomized self-adaptation in evolution strategies. *Evol. Comput.* **9**, 159–195 (2001).
- This article introduced the CMA-ES algorithm, widely regarded as the state of the art in numerical optimization.**
- Bäck, T., Fossette, C. & Krause, P. *Contemporary Evolution Strategies* (Springer, 2013).
- Yao, X. Evolving artificial neural networks. *Proc. IEEE* **87**, 1423–1447 (1999).
- This landmark paper, which was the winner of the 2001 Institute of Electrical and Electronics Engineers Donald G. Fink Prize Paper Award, brought together different strands of research and drew attention to the potential benefits of combining these two forms of learning.**
- Floreano, D., Dürr, P. & Mattiussi, C. Neuroevolution: from architectures to learning. *Evol. Intel.* **1**, 47–62 (2008).
- Barros, R. C., Basgalupp, M. P., de Carvalho, A. C. P. L. F. & Freitas, A. A. A survey of evolutionary algorithms for decision-tree induction. *IEEE Trans. Syst. Man Cybern. C* **42**, 291–312 (2012).
- Widera, P., Garibaldi, J. M. & Krasnogor, N. GP challenge: evolving energy function for protein structure prediction. *Genet. Program. Evolvable Mach.* **11**, 61–88 (2010).
- Filipić, B., Urbančić, T. & Križman, V. A combined machine learning and genetic algorithm approach to controller design. *Eng. Appl. Artif. Intell.* **12**, 401–409 (1999).
- Watson, R. A., Ficici, S. G. & Pollack, J. B. Embodied evolution: distributing an evolutionary algorithm in a population of robots. *Robot. Auton. Syst.* **39**, 1–18 (2002).
- Bredeche, N., Montanier, J. M., Liu, W. & Winfield, A. F. T. Environment-driven distributed evolutionary adaptation in a population of autonomous robotic agents. *Math. Comput. Model. Dyn. Syst.* **18**, 101–129 (2012).
- Nolfi, S. & Floreano, D. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines* (MIT Press, 2000).
- Bongard, J. Evolutionary robotics. *Commun. ACM* **56**, 74–85 (2013).
- Floreano, D. & Keller, L. Evolution of adaptive behavior in robots by means of Darwinian selection. *PLoS Biol.* **8**, e1000292 (2010).
- Hinton, G. E. & Nowlan, S. J. How learning can guide evolution. *Complex Syst.* **1**, 495–502 (1987).
- This seminal paper showed that learning can guide evolution even though characteristics acquired by the phenotype are not communicated to the genotype.**
- Borenstein, E., Meilijson, I. & Ruppin, E. The effect of phenotypic plasticity on evolution in multi-peaked fitness landscapes. *J. Evol. Biol.* **19**, 1555–1570 (2006).
- Paenke, I., Jin, Y. & Branke, J. Balancing population and individual level of adaptation in changing environments. *Adapt. Behav.* **17**, 153–174 (2009).
- Chen, X. S., Ong, Y. S., Lim, M. H. & Tan, K. C. A. Multi-facet survey on memetic computation. *IEEE Trans. Evol. Comput.* **15**, 591–607 (2011).
- Krasnogor, N. & Smith, J. E. A tutorial for competent memetic algorithms: model, taxonomy and design issues. *IEEE Trans. Evol. Comput.* **9**, 474–488 (2005).
- Smith, J. E., Clark, A. R., Stagemeier, A. T. & Serpell, M. C. A genetic approach to statistical disclosure control. *IEEE Trans. Evol. Comput.* **16**, 431–441 (2012).
- Bentley, P. & Corne, D. *Creative Evolutionary Systems* (Morgan Kaufmann, 2002).
- Romero, J. J. & Machado, P. (eds). *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music* (Springer, 2008).
- Secretan, J. et al. Picbreeder: a case study in collaborative evolutionary exploration of design space. *Evol. Comput.* **19**, 373–403 (2011).
- Bentley, P. *Evolutionary Design by Computers* (Morgan Kaufmann, 1999).
- Hingston, P. F., Barone, L. C. & Michalewicz, Z. (eds). *Advances in Evolutionary Design* (Springer, 2008).
- Koza, J. R. Human-competitive results produced by genetic programming. *Genet. Program. Evolvable Mach.* **11**, 251–284 (2010).
- Offers quantifiable definitions for human competitiveness and a well-documented overview of success stories, including the first patents thought to be granted to inventions created by artificial intelligence.**
- Eiben, A. E. & Rudolph, G. Theory of evolutionary algorithms: a bird's eye view. *Theor. Comput. Sci.* **229**, 3–9 (1999).
- Wolpert, D. H. & Macready, W. G. No free lunch theorems for optimisation. *IEEE Trans. Evol. Comput.* **1**, 67–82 (1997).
- This paper reported game-changing results that supported the shift in focus in evolutionary computing and other fields away from the search for a 'super solver', and inspired insightful discussions that are still ongoing.**
- Rudolph, G. Convergence analysis of canonical genetic algorithms. *IEEE Trans. Neural Netw.* **5**, 96–101 (1994).
- Lehre, P. R. & Yao, X. On the impact of mutation-selection balance on the runtime of evolutionary algorithms. *IEEE Trans. Evol. Comput.* **16**, 225–241 (2012).
- Jansen, T. *Analyzing Evolutionary Algorithms: The Computer Science Perspective* (Springer, 2005).
- Borenstein, Y. & Moraglio, A. (eds). *Theory and Principled Methods for Designing Metaheuristics* (Springer, 2014).
- This text provides good coverage of a range of recent approaches and results in the theory of evolutionary algorithms.**
- Eiben, A. E., Hinterding, R. & Michalewicz, Z. Parameter control in evolutionary algorithms. *IEEE Trans. Evol. Comput.* **3**, 124–141 (1999).
- This paper had a long-lasting effect by putting the issue of parameter calibration on the research agenda and establishing the corresponding conceptual framework.**
- Bartz-Beielstein, T. T. *Experimental Research in Evolutionary Computation: the New Experimentalism* (Springer, 2006).
- Hutter, F., Hoos, H. H., Leyton-Brown, K. & Stützle, T. ParamILS: an automatic algorithm configuration framework. *J. Artif. Intell. Res.* **36**, 267–306 (2009).
- Eiben, A. E. & Smit, S. K. Parameter tuning for configuring and analyzing evolutionary algorithms. *Swarm Evol. Comput.* **1**, 19–31 (2011).
- Bartz-Beielstein, T. & Preuss, M. in *Theory and Principled Methods for Designing Metaheuristics* (eds Borenstein, Y. & Moraglio, A.), 205–245 (Springer, 2014).
- Lobo, F. J., Lima, C. F., Michalewicz, Z. (eds). *Parameter Setting in Evolutionary Algorithms* (Springer, 2007).

67. Serpell, M. & Smith, J. E. Self-adaption of mutation operator and probability for permutation representations in genetic algorithms. *Evol. Comput.* **18**, 491–514 (2010).

68. Fialho, A., Da Costa, L., Schoenauer, M. & Sebag, M. Analyzing bandit-based adaptive operator selection mechanisms. *Ann. Math. Artif. Intell.* **60**, 25–64 (2010).

69. Karafotias, G., Hoogendoorn, M. & Eiben, A. E. Parameter control in evolutionary algorithms: trends and challenges. *IEEE Trans. Evol. Comput.* **19**, 167–187 (2015).

This is a recent follow up to ref. 61.

70. Jin, Y. A comprehensive survey of fitness approximation in evolutionary computation. *Soft Comput.* **9**, 3–12 (2005).

71. Jin, Y. Surrogate-assisted evolutionary computation: recent advances and future challenges. *Swarm Evol. Comput.* **1**, 61–70 (2011).

72. Loshchilov, I., Schoenauer, M. & Sebag, M. Self-adaptive surrogate-assisted covariance matrix adaptation evolution strategy. In *Proc. Conference on Genetic and Evolutionary Computation* (eds Soule, T. & Moore, J. H.) 321–328 (ACM, 2012).

73. Zaefferer, M. et al. Efficient global optimization for combinatorial problems. In *Proc. Conference on Genetic and Evolutionary Computation* (eds Igel, C. & Arnold, D. V.) 871–878 (ACM, 2014).

74. Deb, K. *Multi-objective Optimization Using Evolutionary Algorithms* (Wiley, 2001).

75. Zhang, Q. & Li, H. MOEA/D: a multi-objective evolutionary algorithm based on decomposition. *IEEE Trans. Evol. Comput.* **11**, 712–731 (2007).

76. Deb, K. & Jain, H. An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: solving problems with box constraints. *IEEE Trans. Evol. Comput.* **18**, 577–601 (2014).

77. Jain, H. & Deb, K. An evolutionary many-objective optimization algorithm using reference-point based non-dominated sorting approach, part II: handling constraints and extending to an adaptive approach. *IEEE Trans. Evol. Comput.* **18**, 602–622 (2014).

78. Branke, J., Greco, S., Slowinski, R. & Zielniewicz, P. Learning value functions in interactive evolutionary multiobjective optimization. *IEEE Trans. Evol. Comput.* **19**, 88–102 (2015).

79. Stanley, K. O. Compositional pattern producing networks: a novel abstraction of development. *Genet. Program. Evolvable Mach.* **8**, 131–162 (2007).

80. O'Reilly, U.-M. & Hemberg, H. Integrating generative growth and evolutionary computation for form exploration. *Genet. Program. Evolvable Mach.* **8**, 163–186 (2007).

81. Clune, J., Stanley, K. O., Pennock, R. & Ofria, C. On the performance of indirect encoding across the continuum of regularity. *IEEE Trans. Evol. Comput.* **15**, 346–367 (2011).

82. Jin, Y. & Meng, Y. Morphogenetic robotics: an emerging new field in developmental robotics. *IEEE Trans. Syst. Man Cybern. C* **41**, 145–160 (2011).

83. Doursat, R., Sayama, H. & Michel, O. (eds). *Morphogenetic Engineering: Toward Programmable Complex Systems* (Springer, 2013).

84. Doncieux, S., Bredenche, N. & Mouret, J.-B. (eds). *New Horizons in Evolutionary Robotics* (Springer, 2011).

85. Vargas, P. A., Di Paolo, E. A., Harvey, I. & Husbands, P. (eds). *The Horizons of Evolutionary Robotics* (MIT Press, 2014).

86. Harman, M. & McMinn, P. A theoretical and empirical study of search-based testing: local, global, and hybrid search. *IEEE Trans. Softw. Eng.* **36**, 226–247 (2010).

87. Prene, R. & Bull, L. Towards the coevolution of novel vertical-axis wind turbines. *IEEE Trans. Evol. Comput.* **19**, 284–294 (2015).

88. Banzhaf, W. et al. From artificial evolution to computational evolution: a research agenda. *Nature Rev. Genet.* **7**, 729–735 (2006).

89. Maynard Smith, J. Byte-sized evolution. *Nature* **355**, 772–773 (1992).

90. Waibel, M., Floreano, D. & Keller, L. A quantitative test of Hamilton's rule for the evolution of altruism. *PLoS Biol.* **9**, e1000615 (2011).

91. Long, J. *Darwin's Devices: What Evolving Robots Can Teach Us About the History of Life and the Future of Technology* (Basic Books, 2012).

92. Virgo, N., Fernando, C., Bigge, B. & Husbands, P. Evolvable physical self-replicators. *Artif. Life* **18**, 129–142 (2012).

93. Bongard, J. & Lipson, H. Evolved machines shed light on robustness and resilience. *Proc. IEEE* **102**, 899–914 (2014).

94. Bongard, J. Morphological change in machines accelerates the evolution of robust behavior. *Proc. Natl Acad. Sci. USA* **108**, 1234–1239 (2011).

This article demonstrated a hitherto unknown relationship between development, evolution, morphology and the neural control of behaviour, as phrased by the title.

95. Eiben, A. E. Grand challenges for evolutionary robotics. *Front. Robot. AI* **1**, <http://dx.doi.org/10.3389/frobt.2014.00004> (2014).

96. Fernando, C., Kampis, G. & Szathmáry, E. Evolvability of natural and artificial systems. *Procedia Comput. Sci.* **7**, 73–76 (2011).

Acknowledgements The authors would like to thank the editors and reviewers, as well as A. Adamatzky, L. Bull, B. Filipic and M. Schoenauer for providing helpful insights on earlier versions of this manuscript.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The author declares no competing financial interests. Readers are welcome to comment on the online version of this paper at go.nature.com/4r5hj1. Correspondence should be addressed to A.E.E. (a.e.eiben@vu.nl).

New species from Ethiopia further expands Middle Pliocene hominin diversity

Yohannes Haile-Selassie^{1,2}, Luis Gibert³, Stephanie M. Melillo⁴, Timothy M. Ryan⁵, Mulugeta Alene⁶, Alan Deino⁷, Naomi E. Levin⁸, Gary Scott⁷ & Beverly Z. Saylor²

Middle Pliocene hominin species diversity has been a subject of debate over the past two decades, particularly after the naming of *Australopithecus bahrelghazali* and *Kenyanthropus platyops* in addition to the well-known species *Australopithecus afarensis*. Further analyses continue to support the proposal that several hominin species co-existed during this time period. Here we recognize a new hominin species (*Australopithecus deyiremeda* sp. nov.) from 3.3–3.5-million-year-old deposits in the Woranso-Mille study area, central Afar, Ethiopia. The new species from Woranso-Mille shows that there were at least two contemporaneous hominin species living in the Afar region of Ethiopia between 3.3 and 3.5 million years ago, and further confirms early hominin taxonomic diversity in eastern Africa during the Middle Pliocene epoch. The morphology of *Au. deyiremeda* also reinforces concerns related to dentognathic (that is, jaws and teeth) homoplasy in Plio-Pleistocene hominins, and shows that some dentognathic features traditionally associated with *Paranthropus* and *Homo* appeared in the fossil record earlier than previously thought.

Recent fossil hominin discoveries from eastern and central Africa have been described as indicating the existence of multiple locomotor adaptations¹ and taxonomic diversity^{2–5} during the Middle Pliocene. Notably, *Kenyanthropus platyops* from Kenya and *Australopithecus bahrelghazali* from Chad were argued to be morphologically distinct from the contemporaneous *Au. afarensis*^{2–5}. Concerns have been raised about the validity of these two taxa, either because identification was based largely on a single specimen or because the type specimen was heavily distorted⁶. Yet further studies addressing these concerns maintain the initial proposal of taxonomic diversity^{4,5}. More recently, we described a 3.4-million-year (Myr)-old partial foot (BRT-VP-2/73) from the Burtele area of Woranso-Mille, Ethiopia¹. The specimen is contemporaneous with *Au. afarensis*, but demonstrates the existence of a distinct mode of bipedal locomotion. The BRT-VP-2/73 foot combines a grasping hallux with lateral rays that exhibit adaptation to hyperdorsiflexion at the metatarsophalangeal joint, similar to the morphology described for *Ardipithecus ramidus*⁷. This specimen was not assigned to any taxon, pending the recovery of cranial and dentognathic fossils from the same geographic and stratigraphic vicinity. Continued fieldwork in the Burtele area has resulted in the recovery of new dentognathic fossil remains, including partial and complete mandibles, a partial maxilla with dentition, and few isolated teeth (Table 1). Unfortunately, none of them is clearly associated with BRT-VP-2/73 (see Supplementary Note 1 for details).

The new remains described here show clear similarities with *Australopithecus* and lack the diagnostic features of *Ardipithecus*, *Paranthropus*, *Kenyanthropus* and *Homo*. Various aspects of dentognathic morphology distinguish these specimens from the contemporaneous *Au. afarensis* and warrant their assignment to a new species (Fig. 1).

Order Primates Linnaeus, 1758

Suborder Anthropoidea Mivart, 1864

Superfamily Hominoidea Gray, 1825

Genus *Australopithecus* Dart, 1925

***Australopithecus deyiremeda* sp. nov.**

Etymology. From the local Afar language terms *deyi*, meaning close, and *remeda*, meaning relative; thus referring to the species being a close relative of all later hominins.

Holotype. The holotype is BRT-VP-3/1 (Fig. 1a–e), a left maxilla with upper second incisor–upper second molar (I^2 – M^2) found by M. Barao on 4 March 2011. The originals of the holotype and paratypes are housed at the Paleoanthropology Laboratory of the National Museum of Ethiopia, Addis Ababa.

Paratypes. The paratypes are BRT-VP-3/14 (Fig. 1f–h), a complete mandible corpus with the apical halves of left lower incisor–right lower incisor (LI_2 – RI_2) and crowns of right lower premolar–lower molar (RP_3 – M_3), lacking the ascending rami; and WYT-VP-2/10 (Fig. 1i–k), a right edentulous mandible. Both specimens were found by A. Asfaw. Detailed anatomical descriptions, additional images, and measurements are provided in Supplementary Note 2, Extended Data Figs 1 and 2, and Extended Data Tables 1, 2 and 5.

Referred specimen. BRT-VP-3/37 (Fig. 1l–p), a right maxillary fragment with P^4 . This specimen was found only 5 m east of BRT-VP-3/14 and both specimens may belong to the same individual. However, this cannot be determined with confidence owing to the fragmentary nature of BRT-VP-3/37.

Localities. The Burtele (BRT) and Waytaleya (WYT) collection areas are located south of the perennial Mille River and north of the Mille-Chifra road. All localities are situated within an area of 3 km² (Fig. 2). The coordinates of the holotype are 11° 27' 43.9'' N and 40° 31' 41.0'' E (GPS datum WGS84).

Provenience. The holotype and paratypes from BRT-VP-3 are surface finds, weathered out of a sandstone at the bottom of a 6-m thick section of sandstone and siltstone capped by a 2-m thick basaltic flow. A tuff radiometrically dated to 3.469 ± 0.008 (mean \pm s.d.) Myr is located approximately 1 m above the basaltic flow. The paratype from WYT-VP-2 was also a surface find, collected from the top of a thin sandstone layer about 15 m above the 3.469-Myr-old tuff. The combined evidence from radiometric, palaeomagnetic and

¹Cleveland Museum of Natural History, Cleveland, Ohio 44106, USA. ²Case Western Reserve University, Cleveland, Ohio 44106, USA. ³University of Barcelona, Martí Franquès s/n, 08028 Barcelona, Spain. ⁴Max Plank Institute for Evolutionary Anthropology, Deutscher Platz 6, D-04103 Leipzig, Germany. ⁵Pennsylvania State University, University Park, Pennsylvania 16802, USA. ⁶Addis Ababa University, PO Box 1176 Addis Ababa, Ethiopia. ⁷Berkeley Geochronology Center, 2455 Ridge Road, Berkeley, California 94709, USA. ⁸Johns Hopkins University, Baltimore, Maryland 21218, USA.

Table 1 | Fossil hominins from BRT and WYT localities

Specimen number	Preserved elements	Date discovered	Discoverer	Measurements
BRT-VP-1/1	Left M_3 crown fragment	20 February 2006	K. Helem	N/A
BRT-VP-1/2	Left P_4	20 February 2006	K. Helem	10.4 (13.0)
BRT-VP-1/13	Right C_1	20 February 2006	C. Mesfin	9.9, 9.3
BRT-VP-1/18	Left M^2 mesial half	20 February 2006	K. Kayranto	N/A, 13.8
BRT-VP-1/118	Left M_3 fragment	17 February 2012	G. Alemayehu	N/A
BRT-VP-1/120	Left P_4	17 February 2012	A. Elema	9.9, 10.1
BRT-VP-2/10	Frontal fragment	20 February 2006	Y. Haile-Selassie	N/A
BRT-VP-2/73	Right partial foot	15 February 2009	S. Melillo	See ref. ¹
BRT-VP-2/89	Right P^4	17 February 2012	H. Gebreyesus	9.6, 12.9
BRT-VP-2/104	Molar fragment	26 February 2013	M. Mekomeli	N/A
BRT-VP-3/1	Left maxilla (I^2-M^2)	4 March 2011	M. Barao	I^2 , 6.9, 6.6; LC^1 , (8.3), (9.4); LP^3 , 7.7, 12.0; LP^4 , 7.6, 11.8; LM^1 , 10.6, 11.0; LM^2 , 11.8, 13.2
BRT-VP-3/14	Mandible with dentition	4 March 2011	A. Asfaw	LI_1 , 5.4, 8.0; LI_2 , 6.5, 9.0; LC_1 , 8.0, 10.7; LP_3 , 7.6, 10.8; LP_4 , 8.2, 10.6; LM_1 , (12.6), (12.6); LM_2 , (14.8), (13.8); LM_3 , (15.8), (13.4)
BRT-VP-3/37	Right maxilla fragment (P^4)	1 March 2013	A. Asfaw	8.2 (12.1)
WYT-VP-2/10	Partial edentulous mandible	5 March 2011	A. Asfaw	N/A

Dental measurements are standard mesiodistal and buccolingual/labiolingual dimensions. Values in parentheses are corrected for breakage or interproximal attrition. All measurements are in millimetres and were taken on the original specimens by the authors. N/A, non-dental specimens and fragments with no measurement.

depositional rate analyses yields estimated minimum and maximum ages of approximately 3.3 and 3.5 Myr, respectively (Fig. 3, Extended Data Fig. 3 and Supplementary Note 3 for details).

Differential diagnosis. *Australopithecus deyiremeda* differs from *Ar. ramidus* by having relatively and absolutely thicker enamel on its molars; a P_4 with three-roots; a robust mandible. It differs from *Au. anamensis* in lacking an extremely receding mandibular symphysis; having a more robust mandibular corpus and a bicuspid P_3 . It is distinguished from *Au. afarensis* by its overall mandibular architecture, including the lack of a lateral corpus hollow; low and anteriorly

positioned origin of the ascending ramus with pronounced takeoff at the P_4/M_1 level. It also differs in features of the maxilla, including an anteriorly positioned zygomatic origin and in dental size and morphology, specifically an absolutely small upper postcanine dentition and an upper canine that is mesiodistally short and lacks lingual relief. *Australopithecus deyiremeda* differs from *Paranthropus* in lacking molarized premolars and reduced incisors. It is distinguished from *K. platyops* by its transversely convex subnasal region, with incisor alveoli considerably anterior to the bi-canine line; I^1 root larger and more rounded than I^2 root; more prognathic nasoalveolar clivus and

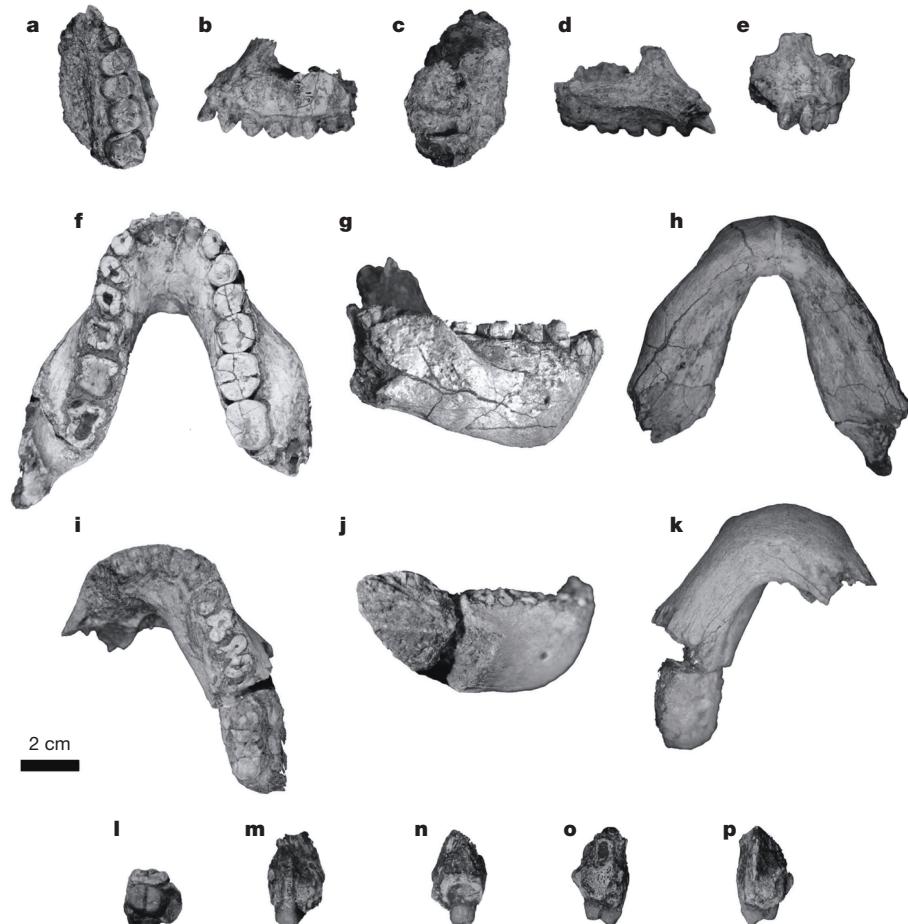


Figure 1 | Holotype BRT-VP-3/1. **a**, Occlusal view. **b**, Lateral view. **c**, Superior view. **d**, Medial view. **e**, Anterior view. Paratype BRT-VP-3/14. **f**, Occlusal view. **g**, Right lateral view. **h**, Basal view. Paratype WYT-VP-2/10. **i**, Occlusal view. **j**, Right lateral view. **k**, Basal view. Referred specimen BRT-VP-3/37. **l**, Occlusal view. **m**, Buccal view. **n**, Lingual view. **o**, Distal view. **p**, Mesial view.

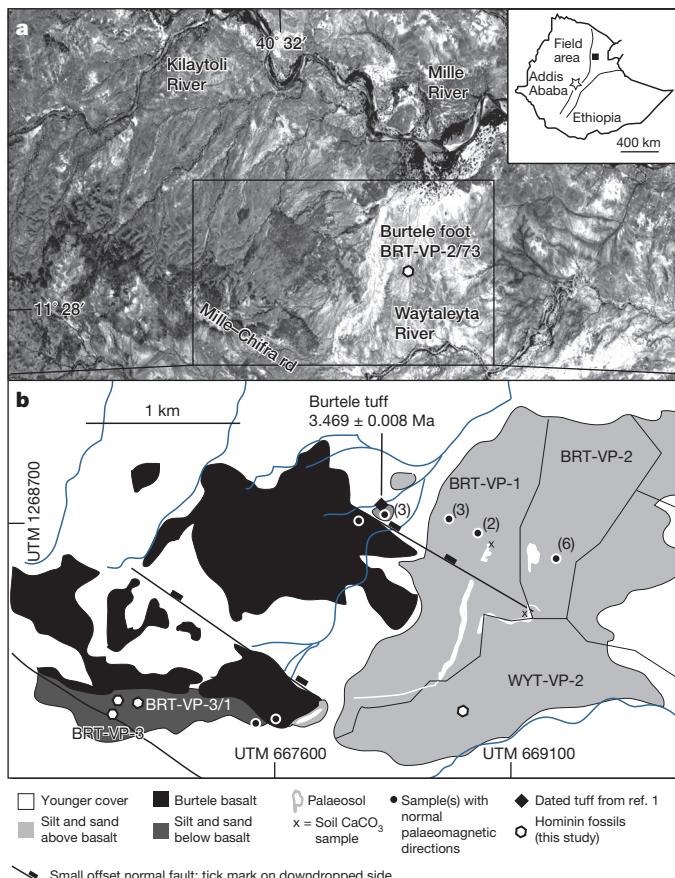


Figure 2 | Location map for the Burtele collection area and the BRT-VP-1, BRT-VP-2, BRT-VP-3 and WYT-VP-2 localities. **a**, Aerial photograph showing the location of the collection area in relation to nearby roads and rivers. **b**, Geological map of the study area showing the outcrop areas of sedimentary and volcanic units and the locations of fossil sites, palaeomagnetic samples (with number of samples indicated in parentheses), the dated Burtele tuff, and location of the holotype and paratypes of *Australopithecus deyiremeda*. The UTM (Universal Transverse Mercator) zone is 37P. The aerial photograph in **a** is reproduced with permission from the Ethiopian Mapping Agency.

I^2 root obliquely implanted, indicating some degree of procumbency; prominent upper canine and P^3 jugae; presence of incipient canine fossa and weakly expressed zygomaticoalveolar crest. *Australopithecus deyiremeda* can be distinguished from *Au. garhi* by its reduced subnasal prognathism and absolutely smaller canines and postcanine dentition. It differs from early *Homo* by its less incisiform canines, asymmetric P_3 , anterior margin of the mandibular ramus arising more anteriorly (opposite distal P_4), mandibular corpus height tapering posteriorly from P_3 to M_3 , and by its anterosuperiorly opening mental foramen.

Comparisons

The holotype maxilla (BRT-VP-3/1) has a more anteriorly positioned zygomatic root compared to *Au. anamensis* and most of the known *Au. afarensis* maxillae. However, it overlaps with most early hominins in the degree of subnasal prognathism (as determined by the subnasal angle) and palate depth^{3,4,8} (see Extended Data Table 2). Although the position of prosthion on BRT-VP-3/1 cannot be determined with certainty, the maximum estimated subnasal angle of this specimen is $\sim 39^\circ$, falling within the range of most early hominins but considerably less than the angle reported for KNM-WT 40000, the holotype of *K. platyops* (47°)⁴.

The incisive foramen in BRT-VP-3/1 is positioned at the distal bicanine line, whereas in *Au. afarensis* maxillae of comparable size (for

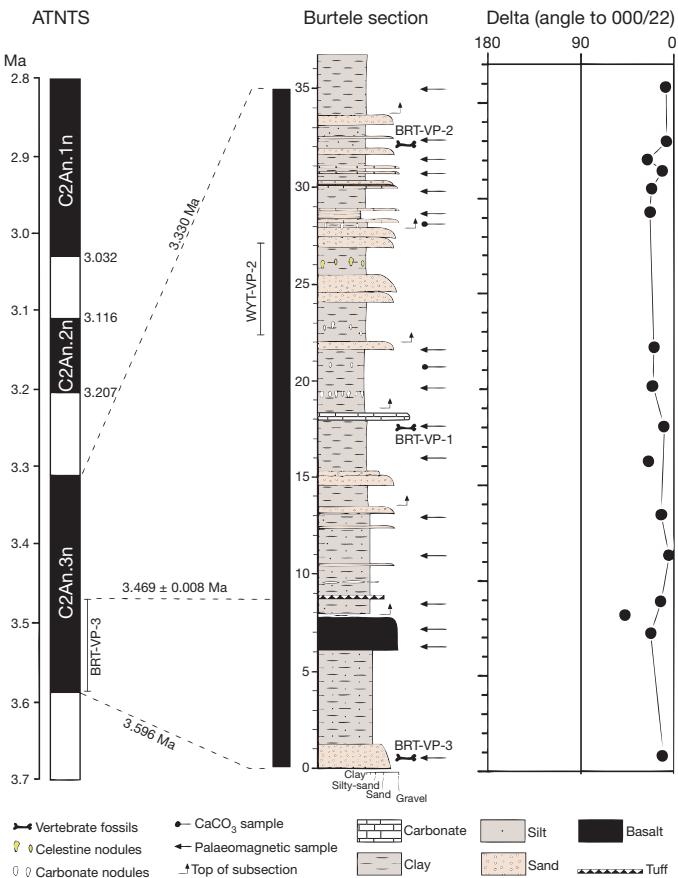


Figure 3 | Palaeomagnetic directions for 16 samples collected from the composite stratigraphic section at Burtele. Directions are shown as the angular distance (δ) to the expected normal direction ($000^\circ, 22^\circ$) for the given latitude. The stratigraphic position of the fossil localities, palaeomagnetic samples and carbonate samples are indicated. All palaeomagnetic samples exhibit normal remanence directions and fall within a single normal magnetozone. K-feldspar phenocrysts from a tuff ~ 1 m above the basalt have been dated by the $^{40}\text{Ar}/^{39}\text{Ar}$ method, yielding a mean age of 3.469 ± 0.008 Myr old¹, indicating that this 36-m section corresponds to the normal magnetozone C2An.3n (3.596–3.330 Ma). A minimum sedimentation accumulation rate of 18.6 cm per 1,000 years is calculated for 26 m of strata deposited in the time period <3.33 –3.47 Ma. Applying this rate, the fossil locality BRT-VP-3 situated 7 m stratigraphically below the tuff (excluding the 2 m basalt, see Supplementary Note 3) has an estimated age of 3.50–3.47 Myr.

example, A.L. 199-1), and early *Homo* specimens such as A.L. 666-1, the foramen is positioned posterior to this line, or at the P^3 level⁸. The palate anterior to the incisive foramen is also flexed inferiorly (see Extended Data Fig. 1i, k, l), a feature considered as derived although it is reported to be variable in *Au. afarensis*⁹. BRT-VP-3/1 shows very little sagittal overlap between the hard palate and the nasoalveolar clivus (3.5 mm) unlike most *Au. afarensis* maxillae (see Extended Data Fig. 4).

Further differences between *Au. deyiremeda* and *Au. afarensis* relate to the morphology of the upper dentition. The cheek teeth of *Au. afarensis* tend to exhibit roughly vertical buccal and lingual walls¹⁰, whereas the buccal and lingual walls of the *Au. deyiremeda* premolars and molars converge occlusally, making the occlusal surface narrower than the base of the crown. The P^3 and P^4 of BRT-VP-3/1 have three roots (Extended Data Fig. 1h, j). This is commonly seen in *Paranthropus* species, but highly variable in *Australopithecus*. *Kenyanthropus platyops* also has three-rooted upper premolars although P^3 – P^4 crown morphology is currently unknown for this species³.

The canine crown of BRT-VP-3/1 is absolutely smaller than almost all known *Australopithecus* upper canines, specifically in its

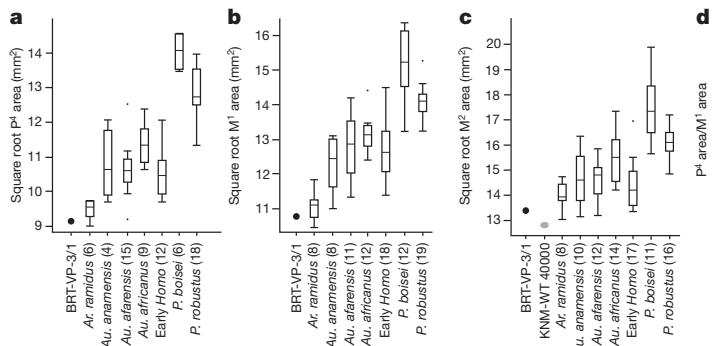


Figure 4 | Box-and-whisker diagrams comparing dental dimensions of the holotype BRT-VP-3/1 with other early hominins. Square roots of P⁴–M² areas (buccolingual \times mesiodistal dimensions) are plotted. The horizontal edges of each box indicate the first quartile (Q_1), the median (Q_2), and the third quartile (Q_3). The superiormost and inferiormost horizontal lines show the maximum and minimum values, respectively, excluding outliers, which are shown above or below these horizontal lines. Metric data for the comparative samples were compiled from refs 9 and 25, 28–33.

mesiodistal dimension, which falls close to the mean of *P. robustus* and *P. boisei* (Extended Data Table 2). It also appears to be long-rooted relative to its crown height, as in *Au. anamensis*^{11–13}. Furthermore, the canine of BRT-VP-3/1 does not have pronounced lingual relief and entirely lacks the prominent mesial and distal vertical lingual grooves commonly seen in *Au. afarensis* upper canines¹⁰. Its distal tubercle is also small and the mesial crown shoulder relatively short compared to the crown height and length of the mesial crest, thus it is less incisiform than *Au. afarensis* upper canines. It retains a more primitive, *Au. anamensis*-like condition in this regard.

The postcanine teeth of BRT-VP-3/1, particularly the P⁴ and M¹, are very small (Fig. 4; see Extended Data Table 2 for summary statistics). The M¹ area of BRT-VP-3/1 is the smallest of all known Pliocene hominins (less *Ar. ramidus*), whereas the P⁴ size is only approximated by A.L. 199-1, the smallest outlier in the *Au. afarensis* hypodigm. The M² also falls at the lower end of the *Au. afarensis* range and is slightly larger than that of KNM-WT 40000, the holotype of *K. platyops*. The P⁴ to M¹ area ratio of BRT-VP-3/1 falls within the range of most Pliocene hominin taxa, but outside the range of *Paranthropus*. In terms of enamel thickness, the two-dimensional linear enamel thickness for the M¹ of BRT-VP-3/1 is similar to that of other early hominins such as *Au. afarensis*, *P. robustus* and *Au. africanus*. However, the two-dimensional and three-dimensional enamel thickness values for the M² of this specimen are high, most similar to *P. robustus* (see Supplementary Note 4 and Extended Data Tables 3 and 4 for details).

Although *Au. deyiremeda* and *K. platyops* show some similarities in maxillary morphology and molar size, they can be distinguished by a suite of qualitative characters related to the morphology of the nasoalveolar clivus, shape of the dental arcade, and morphology of the anterior maxillary dentition. The nasoalveolar clivus in *K. platyops* is flat both sagittally and transversely³, whereas in *Au. deyiremeda* it is anteriorly convex in the transverse plane. In accordance with this difference, *Au. deyiremeda* and *K. platyops* differ in the shape of the anterior dental arcade. On the basis of the position of the preserved I¹ root, the transversely oblique orientation of the I² crown, and its placement relative to the canine and postcanine row, it is clear that the canine and incisors of BRT-VP-3/1 formed an arc with the incisor alveoli well anterior to the bi-canine line, as in *Au. afarensis* specimens such as A.L. 200-1a and A.L. 444-2 (ref. 9; see Extended Data Fig. 1f, g). In *K. platyops*, the incisors are arranged parallel to the bi-canine line as in *Paranthropus* and *Homo rudolfensis*³. *Kenyanthropus platyops* also has a derived low and curved zygomaticoalveolar crest³, whereas this feature is not sufficiently developed to form a visible crest in BRT-VP-3/1. Instead, the lateral alveolar wall blends smoothly with the base of the zygomatic bone. Furthermore, computed tomography scans show that the incisor roots of BRT-VP-3/1 are straight and inclined posterosuperiorly, indicating procumbent incisors (Extended Data Fig. 1i, k). By contrast, the incisor roots of *K. platyops* do not show any sign of procumbency³. On the basis of root cross-section approximately 2 mm below the buccal cervicoenamel line of I², the I¹ root of BRT-VP-3/1 is also rounded and more

robust than the mesiolaterally compressed I² root, unlike KNM-WT 40000 in which the I¹ and I² incisor roots are almost equal in size³ (see Extended Data Fig. 1m). In general, *K. platyops* seems to be more derived than *Au. deyiremeda* in most of its maxillary and dental morphology.

The overall morphology of the paratype mandibles (BRT-VP-3/14 and WYT-VP-2/10; Fig. 1 and Extended Data Fig. 2) is comparable to that of eastern African robust taxa, although their overall size is within the range of *Au. afarensis*. The anterior origin of the ascending ramus at the posterior P₄ level is comparable to *P. boisei* mandibles such as the Peninj mandible and KGA-10-525 (ref. 14), and *P. robustus* specimens such as SK-12. *Australopithecus afarensis* mandibles have the origin of the ascending ramus high and more posterior, mostly distal to the first molar^{9,15}. BRT-VP-3/14 and WYT-VP-2/10 have inflated mandibular corpora, especially at the M₂₋₃ level, and they lack the lateral corpus hollow ubiquitous in *Au. afarensis* mandibles⁹. Corpus dimensions of BRT-VP-3/14 and WYT-VP-2/10 are notable because they are broader for their size than specimens assigned to *Au. afarensis* or early *Homo*. In this respect, they are more similar to *Paranthropus* (Fig. 5 and Supplementary Note 5). Mandibular specimens from Lomekwi whose taxonomic affinities have not been established with certainty (KNM-WT 8556 and KNM-WT 16006)^{3,16}, are described as having robust corpora with an ascending ramus take-off at P₄ and M₁, respectively. However, the robusticity of the Lomekwi specimens is not as pronounced as it is in WYT-VP-2/10 and BRT-VP-3/14. Furthermore, the KNM-WT 8556 premolars are derived: the P₃ has a well-developed metaconid and the P₄ is large and molarized³. By contrast, the P₃ of BRT-VP-3/14 is almost unicuspied, with only a trace of the metaconid, and the P₄ is not molarized.

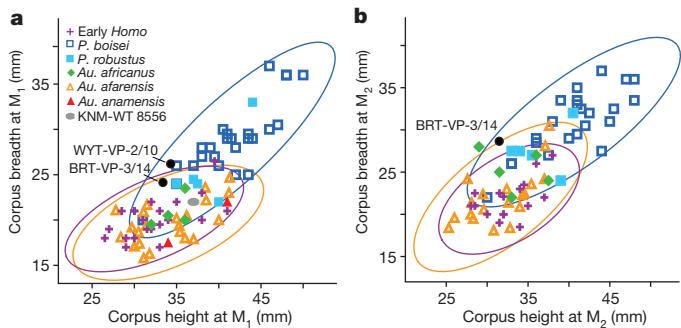


Figure 5 | Bivariate plots of mandibular corpus height and breadth among early hominins. a, Dimensions at the M₁ level. b, Dimensions at the M₂ level. The 95% confidence ellipses are shown for *Au. afarensis*, *P. boisei* and early *Homo* (see Supplementary Note 5 for details). The *Au. deyiremeda* corpora are relatively broad, comparable to *Paranthropus*. WYT-VP-2/10 falls outside the *Au. afarensis* ellipse at the M₁ level. BRT-VP-3/14 (represented by the mean of left and right sides) falls in an area of overlap at the M₁ level but at the edge of the *Au. afarensis* ellipse at the M₂ level. Mandibular measurements for the comparative taxa were taken from the references listed in Fig. 4 and ref. 14.

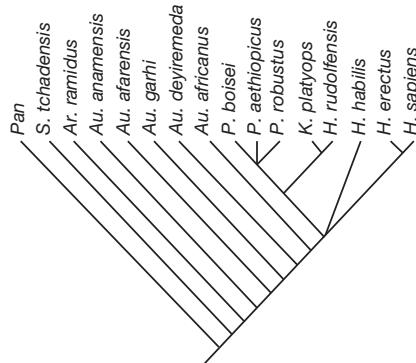


Figure 6 | Cladogram depicting the majority-rule consensus of 17 equally parsimonious trees that result from a phylogenetic analysis of the features preserved in the *Au. deyiremeda* hypodigm. Tree length = 63, consistency index = 0.63. The position of *Au. garhi* and topologies within the *Paranthropus* + *Kenyanthropus* + *Homo* clade are unstable when these features are considered. However, *Au. deyiremeda* is consistently placed as a sister taxon to a clade that includes *Au. africanus*, *Paranthropus*, and *Homo* (see Supplementary Notes 6–8 for further details).

The symphysis of BRT-VP-3/14 is thinner, superoinferiorly longer, and has a more receding external contour than that of WYT-VP-2/10 (Extended Data Fig. 2h, n). In the latter specimen, the superior half of the external contour is almost vertical but the inferior half is extremely receding similar to L.H.-4, the holotype of *Au. afarensis*^{9,17,18}. Further similarities between WYT-VP-2/10 and L.H.-4 include the prominence of the superior and inferior transverse tori (hence the deep genioglossal fossa) and the posterior extension of the inferior transverse torus exceeding that of the superior transverse torus. BRT-VP-3/14 has weak transverse tori and a very shallow genioglossal fossa (similar to *Au. afarensis* specimens such as A.L. 333w-12 and A.L. 315-22) and a pronounced postincisive planum (comparable to A.L. 400-1). These differences are within the range of variation seen in *Au. afarensis* mandibles⁹ and further show intra-taxon variability in mandibular symphyseal morphology. Although it was previously argued that two *Au. bahrelghazali* mandibles (KT 12 and KT 40) have symphyseal morphology that distinguishes them from *Au. afarensis* mandibles⁵, they appear to also fall within the range of *Au. afarensis* variation^{9,19,20}.

In terms of the dentition, BRT-VP-3/14 has small premolars compared to *Au. afarensis* (see Extended Data Table 5 and Supplementary Note 5 for details). The small mandibular premolars appear to conform well with the small upper premolars and M¹ of BRT-VP-3/1. Despite the robust appearance of the mandible, the BRT-VP-3/1 maxilla is only slightly smaller when occluded with the BRT-VP-3/14 mandible (see Extended Data Fig. 5).

The 2.8-Myr-old early *Homo* mandible (LD 350-1) recently reported from Ledi-Geraru²¹ is similar to *Au. deyiremeda* and some *Au. afarensis* specimens in its anterior corpus morphology. However, the derived features that are present in LD 350-1 and diagnostic of *Homo* are absent in the *Au. deyiremeda* mandibles (for example, posteriorly opening mental foramen, comparable anterior and posterior corpus height, M₃ mesiodistally shorter than M₂). Features of *Au. deyiremeda* that are derived relative to *Au. afarensis* and earlier hominins (for example, relatively robust mandibular corpus, anteriorly positioned root of the ascending ramus) are absent in LD 350-1 and other early *Homo* mandibles. Dentally, while the P₄ of BRT-VP-3/14 and LD 350-1 are comparable in size, the M₂ and M₃ of BRT-VP-3/14 are much larger (Extended Data Table 5).

Phylogenetic implications

The taxonomy and phylogenetic relationships among early hominins are becoming more complicated as new taxa are added to the Pliocene fossil record^{2,3,22,23} and the temporal range and systematics of early

Homo are reconsidered^{21,24}. A phylogenetic analysis of the preserved morphological features of *Au. deyiremeda* fails to produce a single most-parsimonious cladogram. However, the results are fairly consistent in placing this species as a sister taxon to a clade including *Au. africanus*, *Paranthropus* and *Homo* (Fig. 6 and Supplementary Notes 6–8). There are many alternative phylogenies consistent with this topology, but the fact that *Au. deyiremeda* is positioned outside the *Paranthropus* and *Homo* clade implies that some features associated with one or both of these taxa are homoplastic. Therefore, the addition of *Au. deyiremeda* reinforces questions about dentognathic homoplasy in later Pliocene and early Pleistocene hominins.

There is now incontrovertible evidence to show that multiple hominins existed contemporaneously in eastern Africa during the Middle Pliocene^{1,3}. Importantly, Woranso-Mille is geographically very close to Hadar (only 35 km north), where *Au. afarensis* is well documented. In combination, this suggests that multiple hominin species overlapped temporally and lived in close geographic proximity. What remains intriguing, and requires further investigation, is how these taxa are related to each other and to later hominins, and what environmental and ecological factors triggered such diversity. This has important implications for resource use and niche partitioning in Pliocene hominin palaeoecology, which can only be understood with the recovery and analysis of more hominin fossils and their associated fauna from Woranso-Mille and other contemporaneous sites.

Received 8 October 2014; accepted 9 April 2015.

- Haile-Selassie, Y. et al. A new hominin foot from Ethiopia shows multiple Pliocene bipedal adaptations. *Nature* **483**, 565–570 (2012).
- Brunet, M. et al. *Australopithecus bahrelghazali*, une nouvelle espèce d'Hominidé ancien de la région de Koro Toro (Tchad). *C. R. Acad. Sci.* **322**, 907–913 (1996).
- Leakey, M. G. et al. New hominin genus from eastern Africa shows diverse middle Pliocene lineages. *Nature* **410**, 433–440 (2001).
- Spoor, F., Leakey, M. G. & Leakey, L. N. Hominin diversity in the middle Pliocene of eastern Africa: the maxilla of KNM-WT 4000. *Phil. Trans. R. Soc. B* **365**, 3377–3388 (2010).
- Guy, F. et al. Symphyseal shape variation in extant and fossil hominoids, and the symphysis of *Australopithecus bahrelghazali*. *J. Hum. Evol.* **55**, 37–47 (2008).
- White, T. Early hominids—diversity or distortion? *Science* **299**, 1994–1997 (2003).
- Lovejoy, C. O., Latimer, B., Suwa, G., Asfaw, B. & White, T. D. Combining prehension and propulsion: the foot of *Ardipithecus ramidus*. *Science* **326**, 72e1–72e8 (2009).
- Kimbel, W. H., Johanson, D. C. & Rak, Y. Systematic assessment of a maxilla of *Homo* from Hadar, Ethiopia. *Am. J. Phys. Anthropol.* **103**, 235–262 (1997).
- Kimbel, W. H., Rak, Y. & Johanson, D. C. *The Skull of Australopithecus afarensis* (Oxford Univ. Press, 2004).
- Johanson, D. C., White, T. D. & Coppens, Y. Dental remains from the Hadar Formation, Ethiopia: 1974–1977 collections. *Am. J. Phys. Anthropol.* **57**, 545–603 (1982).
- Ward, C. V., Leakey, M. G. & Walker, A. Morphology of *Australopithecus anamensis* from Kanapoi and Allia Bay, Kenya. *J. Hum. Evol.* **41**, 255–368 (2001).
- Manthi, F., Plavcan, M. & Ward, C. New hominin fossils from Kanapoi, Kenya: mosaic evolution of canine evolution in early hominins. *So. Afr. J. Sci.* **108**, 1–9 (2012).
- Ward, C. V., Manthi, F. K. & Plavcan, J. M. New fossils of *Australopithecus anamensis* from Kanapoi, West Turkana, Kenya (2003–2008). *J. Hum. Evol.* **65**, 501–524 (2013).
- Suwa, G. et al. The first skull of *Australopithecus boisei*. *Nature* **389**, 489–492 (1997).
- Kimbel, W. H. & Delezenne, L. K. ‘Lucy’ redux: a review of research on *Australopithecus afarensis*. *Am. J. Phys. Anthropol.* **49**, 2–48 (2009).
- Brown, B., Brown, F. H. & Walker, A. New hominids from the Lake Turkana Basin, Kenya. *J. Hum. Evol.* **41**, 29–44 (2001).
- White, T. D. & Johanson, D. C. Pliocene hominid mandibles from the Hadar Formation, Ethiopia: 1974–1977 collections. *Am. J. Phys. Anthropol.* **57**, 501–544 (1982).
- Johanson, D. C., White, T. & Coppens, Y. *A New Species of the Genus Australopithecus (Primates: Hominidae) From the Pliocene of Eastern Africa* (Kirtlandia) (Cleveland Museum of Natural History, 1978).
- White, T. D., Suwa, G., Simpson, S. W. & Asfaw, B. Jaws and teeth of *Australopithecus afarensis* from Maka, Middle Awash, Ethiopia. *Am. J. Phys. Anthropol.* **111**, 45–68 (2000).
- Wood, B. & Richmond, B. Human evolution: taxonomy and paleobiology. *J. Anat.* **197**, 19–60 (2000).
- Villmoare, B. et al. Early *Homo* at 2.8 Ma from Ledi-Geraru, Afar, Ethiopia. *Science* **347**, 1352–1355 (2015).
- Asfaw, B. et al. *Australopithecus garhi*: a new species of early hominid from Ethiopia. *Science* **284**, 629–635 (1999).
- Leakey, M. G., Feibel, C. S., McDougall, I. & Walker, A. New four-million-year-old hominid species from Kanapoi and Allia Bay, Kenya. *Nature* **376**, 565–571 (1995).

24. Spoor, F. et al. Reconstructed *Homo habilis* type OH 7 suggests deep-rooted species diversity in early *Homo*. *Nature* **519**, 83–86 (2015).
25. Wood, B. *Koobi Fora Research Project Vol. 4* (Clarendon, 1991).
26. White, T. D., Johanson, D. C. & Kimbel, W. H. in *New Interpretations of Ape and Human Ancestry* (eds Ciochon, R. L. & Corruccini, R. S.) 721–780 (Plenum, 1983).
27. Harrison, T. in *Paleontology and Geology of Laetoli: Human Evolution in Context* (ed. Harrison, T.) Vol. 2 141–188 (Dordrecht: Springer, 2011).
28. White, T. D. Additional fossil hominids from Laetoli, Tanzania: 1976–1979 specimens. *Am. J. Phys. Anthropol.* **53**, 487–504 (1980).
29. White, T. D. New fossil hominids from Laetoli, Tanzania. *Am. J. Phys. Anthropol.* **46**, 197–229 (1977).
30. White, T. D. et al. Asa Issie, Aramis and the origin of *Australopithecus*. *Nature* **440**, 883–889 (2006).
31. White, T. D., Suwa, G., Simpson, S. W. & Asfaw, B. Jaws and teeth of *Australopithecus afarensis* from Maka, Middle Awash, Ethiopia. *Am. J. Phys. Anthropol.* **111**, 45–68 (2000).
32. Suwa, G. et al. Paleobiological implications of the *Ardipithecus ramidus* dentition. *Science* **326**, 94–99 (2009).
33. Leakey, M. G. et al. New fossils from Koobi Fora in northern Kenya confirm taxonomic diversity in early *Homo*. *Nature* **488**, 201–204 (2012).
34. Olejniczak, A. J. et al. Three-dimensional molar enamel distribution and thickness in *Australopithecus* and *Paranthropus*. *Biol. Lett.* **4**, 406–410 (2008).

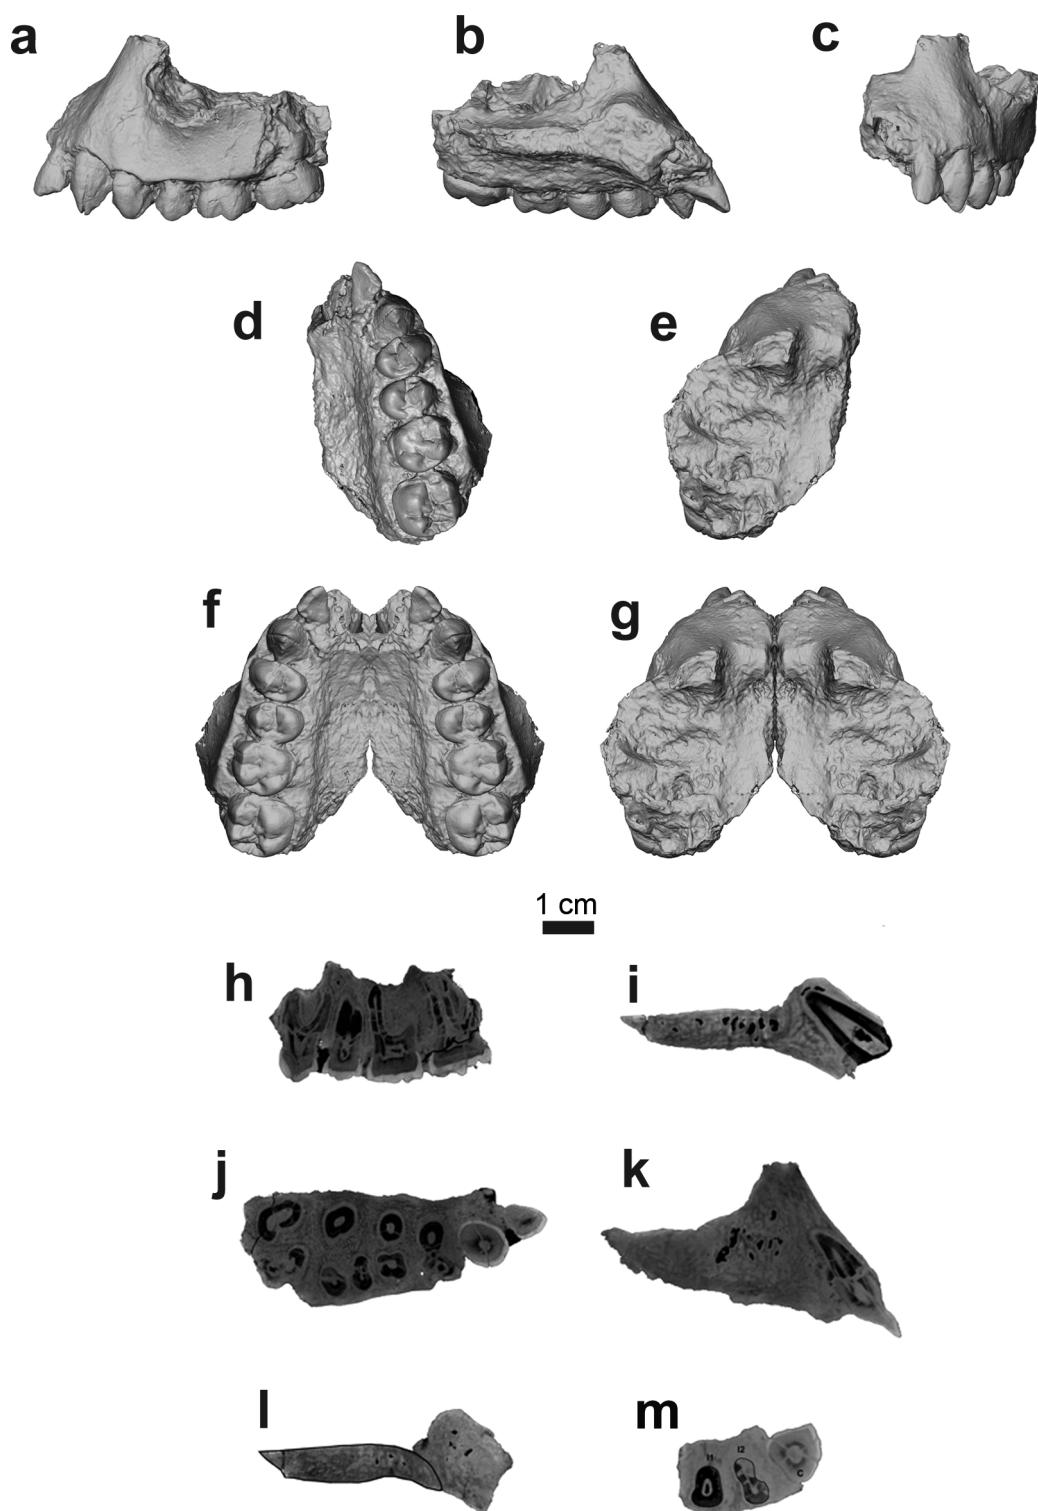
Supplementary Information is available in the online version of the paper.

Acknowledgements We thank the Authority for Research and Conservation of Cultural Heritage and the Afar Regional State of Ethiopia for permission to conduct field and

laboratory research, the Afar people of the Woranso–Mille area for their hospitality, and the project's fieldwork crew members for their incredible contributions in our fieldwork endeavours. We would also like to thank S. W. Simpson and W. H. Kimbel for their constructive comments and discussions throughout the preparation of this manuscript; G. Suwa, T. White and B. Asfaw for access to the original *Ardipithecus ramidus* material; W. H. Kimbel for access to the original *Australopithecus afarensis* material; F. Spoor for discussion of *Kenyanthropus platyops* morphology and Lomekwi mandibular metrics; D. F. Su for comments and assistance with the figures. This research was financially supported by grants from the LSB Leakey Foundation, the National Geographic Society, the Cleveland Museum of Natural History, and the National Science Foundation (BCS-0234320, BCS-0321893, BCS-0542037, BCS-1124705, BCS-1124713, BCS-1124716, BCS-1125157 and BCS-1125345).

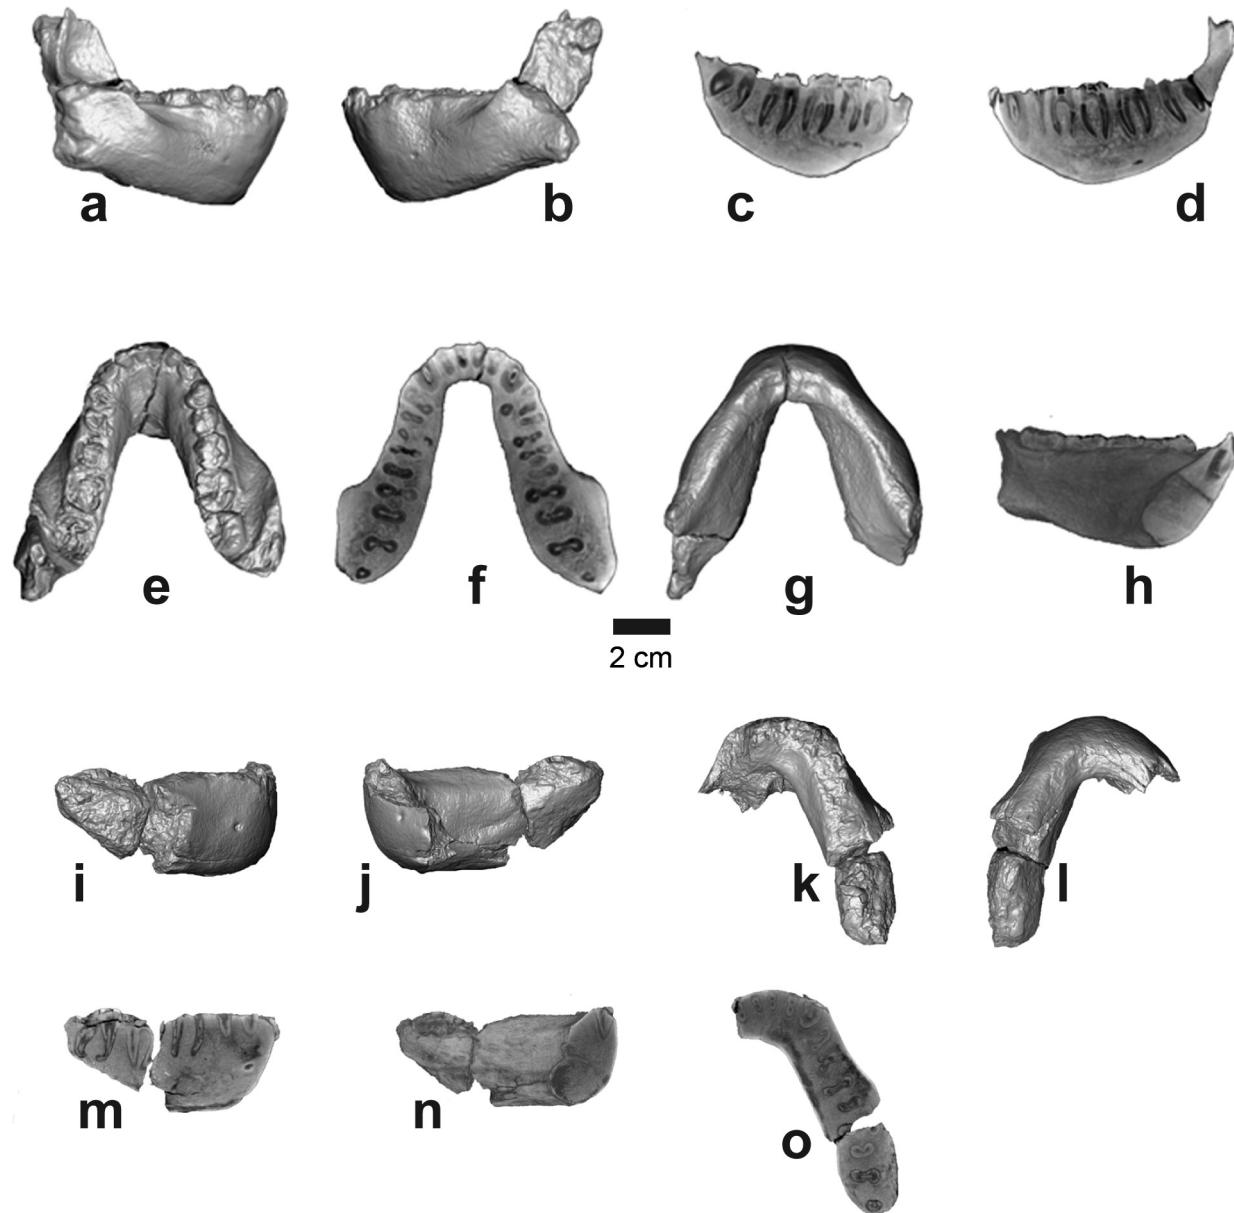
Author Contributions Y.H.-S. directed the field research. A.D., B.Z.S., M.A., L.G. and S.M.M. participated in the field research. T.M.R. conducted the computed tomography scanning and made the virtual reconstructions. A.D., B.Z.S., L.G. and M.A. studied the geological context, while L.G. and G.S. conducted the palaeomagnetic analysis. S.M.M. conducted the phylogenetic analysis. Y.H.-S. and S.M.M. made comparative observations and carried out analyses. Y.H.-S. took the lead in writing the paper with contributions from all coauthors.

Author Information The LSID urn:lsid:zoobank.org:pub:0C492889-01AC-4CDD-96FD-0E08A51F3CBB has been deposited in ZooBank. Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to Y.H.-S. (yhailese@cmnh.org).



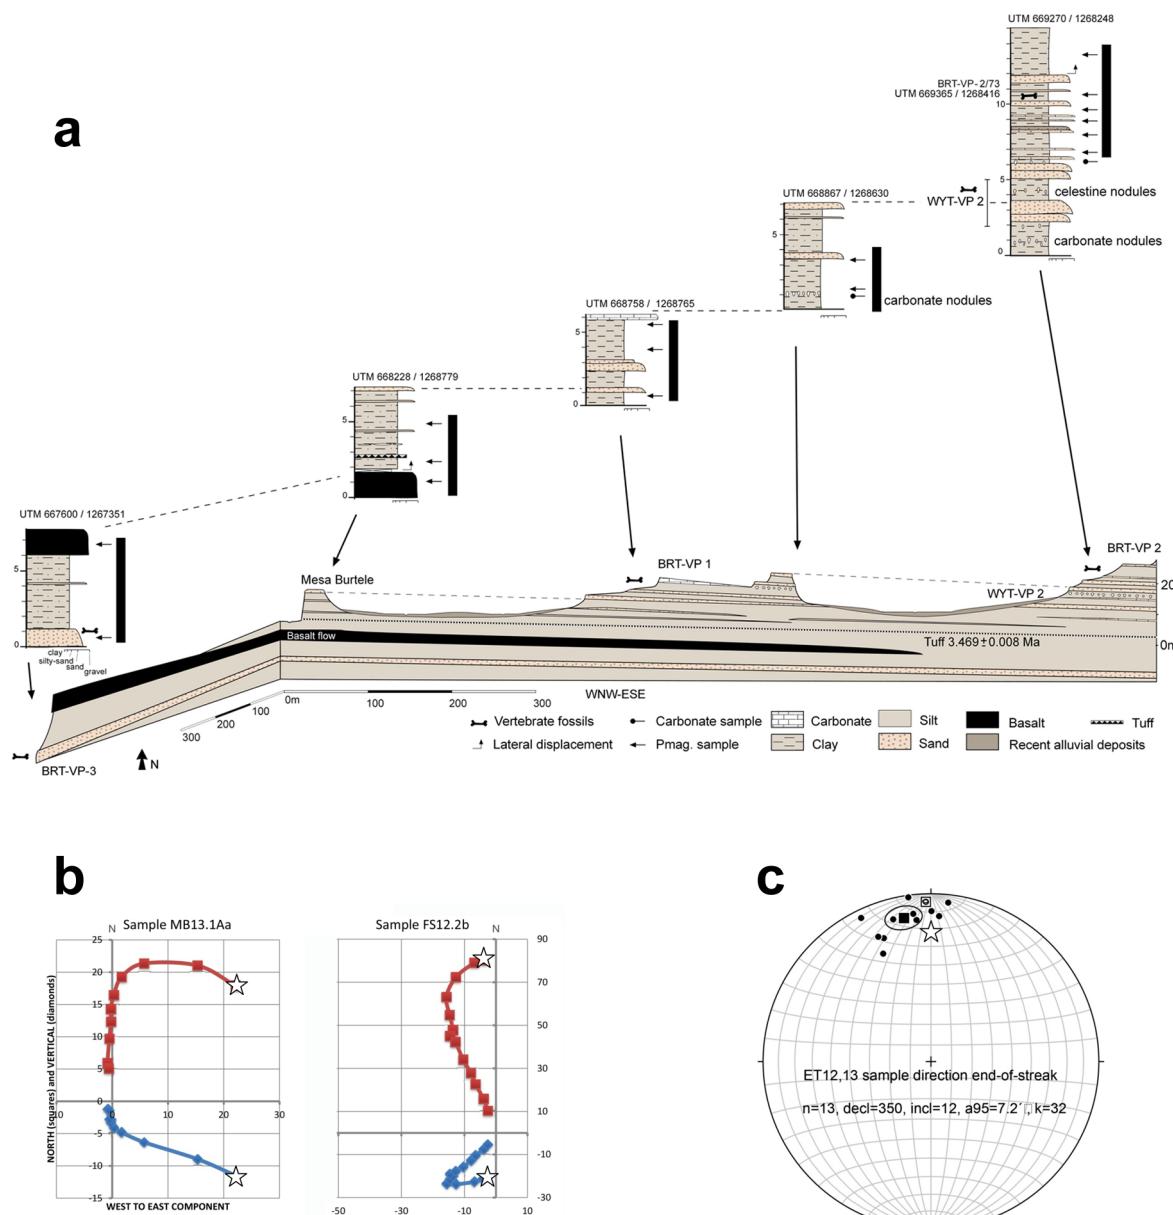
Extended Data Figure 1 | Computed-tomography-based visualization of BRT-VP-3/1. **a**, Lateral view. **b**, Medial view. **c**, Anterior view. **d**, Palatal view. **e**, Superior view. **f**, Palatal view mirror-imaged on midline. **g**, Superior view mirror-imaged on midline. Computed-tomography-based cross-sections of BRT-VP-3/1. **h**, Sagittal cross-section along the centre of the dental row

showing root morphology. **i**, Sagittal cross-section at I^1 . **j**, Transverse cross-section along the roots showing the number of roots of each tooth. **k**, Sagittal cross-section at I^2 . **l**, Midsagittal cross-section showing the shape of the palatine process and nasoalveolar clivus. **m**, Transverse cross-section across the incisors and the canine.



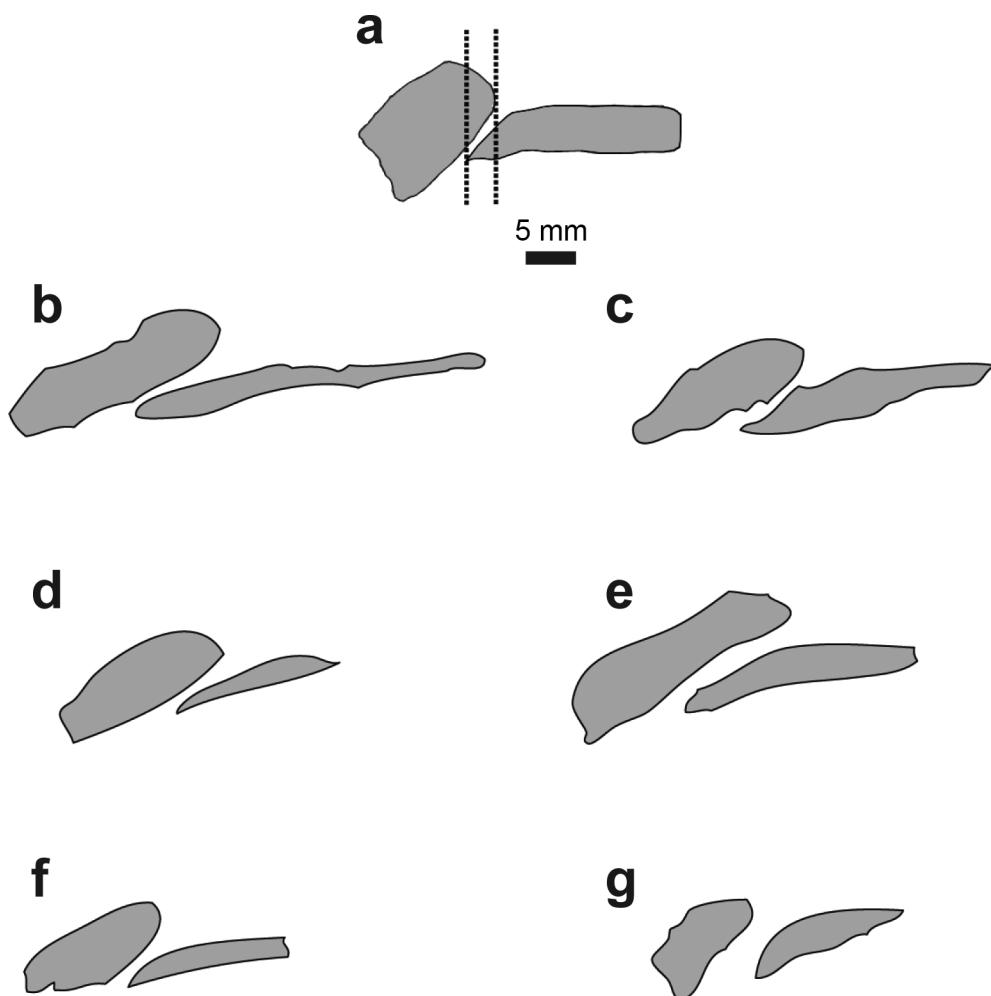
Extended Data Figure 2 | Computed-tomography-based visualization of BRT-VP-3/14. **a**, Right lateral view. **b**, Left lateral view. **c**, Sagittal cross-section along the centre of the right dental row showing root morphology. **d**, Sagittal cross-section along the centre of the left dental row showing root morphology. **e**, Occlusal view. **f**, Transverse cross-section along the roots showing the number of roots of each tooth. Note that the premolars have several roots.

g, Basal view. **h**, Symphyseal cross-section. Computed-tomography-based visualization of WYT-VP-2/10. **i**, Right lateral view. **j**, Medial view. **k**, Occlusal view. **l**, Basal view. **m**, Sagittal cross-section along the centre of the dental row showing root morphology. **n**, Symphyseal cross-section. **o**, Transverse cross-section along the roots showing the number of roots of each tooth.



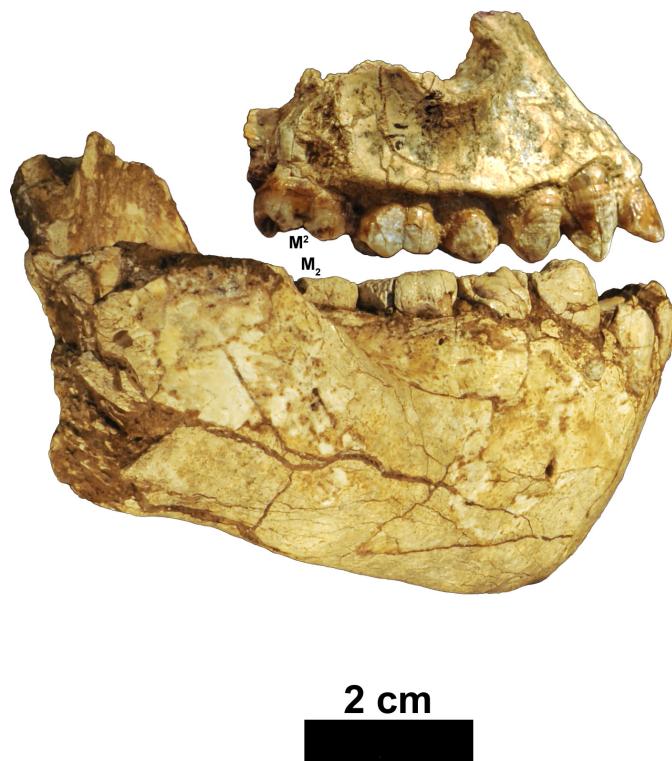
Extended Data Figure 3 | Magnetostratigraphy. **a**, Magnetostratigraphy for the Burtele area, measured in five subsections beginning with the oldest strata in the southwest (BRT-VP-3), to the youngest at BRT-VP-2 (see Fig. 2 for subsections locations). This stratigraphy consists predominantly of claystones/siltstones alternating with sandstones, and includes a 20-cm carbonate bed, a 2–6-m basalt flow, and a 5-cm tuff bed. The sandstones occupy slightly sinuous fluvial channels with a low width/depth ratio in a weakly confined setting. The sandstone beds are continuous and interstratified with pedogenically modified fine-grained overbank and ephemeral lake deposits. The proportion of sandstones increases upwards in the section, where it is accompanied by the occurrence of celestine nodules indicating increased aridity. Sixteen palaeomagnetic samples were obtained from fine-grained lithologies. All

samples demonstrate normal polarity (black bars), interpreted as belonging to a single polarity chron C2An.3n (3.596–3.330 Myr old) of the Astronomically Tuned Neogene Time Scale (ATNTS2004). **b**, Orthogonal demagnetization diagrams showing vector endpoints after alternating field and thermal demagnetization for samples MB13.1Aa and FS12.2b. Horizontal projections are shown as squares, vertical projection as diamonds, and NRM (4 mT) starting point as star. **c**, Palaeomagnetic directions for the Burtele stratigraphy from samples collected along different stratigraphic horizons. Stereographic projection referenced to magnetic North (declination 2° E), with solid symbols on lower hemisphere (plus inclination). Represented directions are the mean directions for each sample from three specimens. The start shows the location of the expected normal direction for this latitude (000/22).



Extended Data Figure 4 | Comparisons of midsagittal cross section of the hard palate and nasoalveolar clivus. a, BRT-VP-3/1. b, A.L. 444-2. c, A.L. 200-1. d, A.L. 486-1. e, A.L. 427-1. f, A.L. 199-1. g, A.L. 442-1. Note that the overlap

between the hard palate and the nasolalveolar clivus of BRT-VP-3/1 is small relative to most *Au. afarensis* specimens. Midsagittal cross-sections of *Au. afarensis* maxillae were modified from Fig. 5.22 in ref. 9.



Extended Data Figure 5 | BRT-VP-3/1 (reversed) and BRT-VP-3/14 shown in occlusion. The upper canine is aligned mesial to the P_3 . Despite the apparently large size of the mandible (BRT-VP-3/14), the maxilla (BRT-VP-3/

1) is only slightly smaller. The position of the second molars is indicated to show that the canine-to-second-molar length is comparable in both specimens.

Extended Data Table 1 | Measurements of the paratype mandibles

Measurement	BRT-VP-3/14	WYT-VP-2/10
1 Symphyseal height	42.2	40.7
2 Maximum symphyseal depth	19.9	23.3
3 Horizontal position of mental foramen	P3/P4	P3
4 Mental foramen height from base	19.4	18
5 Mental foramen height from alveolar margin	~17.4	~16.6
6 Corpus height at P ₄	36.6, 35.8	35.7
7 Corpus width at P ₄	21.1, 21.8	23.8
8 Corpus height at P ₄ /M ₁	36.3, 35.9	34.8
9 Corpus width at P ₄ /M ₁	22.8, 22.7	27.11
10 Corpus height at M ₁	33.8, 33.0	34.3
11 Corpus width at M ₁	24.5, 23.6	26.2
12 Corpus height at M ₁ /M ₂	33.3, 31.9	-
13 Corpus width at M ₁ /M ₂	25.3, 28.0	-
14 Corpus height at M ₂	32.4, 30.5	-
15 Corpus width at M ₂	28.2, 29	-
16 Corpus height at M ₂ /M ₃	30.8, 28.5	-
17 Corpus width at M ₂ /M ₃	28.1, 29.5	-
18 P ₄ -P ₄ distance at mid-crown on lingual side	29.9	-
19 M ₁ -M ₁ distance at mid-crown on lingual side	~33.5	-

Corpus dimensions of BRT-VP-3/14 are reported for both right and left sides. Measurements 1–5 were taken following the methods used in ref. 25. Measurements 6–19 were taken following the methods used in ref. 17. All measurements are in millimetres.

Extended Data Table 2 | Statistical summary of measurements of early hominin maxillae and upper dentition

		Subnasal Angle	Palate Depth	C ¹ MD	$\sqrt{P^4}$ Area	$\sqrt{M^1}$ Area	$\sqrt{M^2}$ Area	$\sqrt{P^4}$ Area / $\sqrt{M^1}$ Area
BRT-VP-3/1		~39	~13.4	8.3	9.1	10.8	12.5	0.84
<i>Au. anamensis</i>	n	1	—	2	4	8	10	4
	mean	27	—	10.8	10.6	12.5	13.7	0.90
	range	—	—	9.9 - 11.7	9.7 - 12.1	11.0 - 13.1	12.3 - 15.5	0.79 - 0.92
	s.d.	—	—	—	1.0	0.74	1.02	0.06
<i>Au. afarensis</i>	n	6	6	10	15	11	12	11
	mean	34.6	11.3	9.9	10.6	12.9	13.9	0.81
	range	29 - 39	8.5 - 14	8.9 - 11.6	9.2 - 12.5	11.4 - 14.2	12.3 - 14.9	0.74 - 0.97
	s.d.	3.5	1.8	0.74	0.72	0.88	0.73	0.07
<i>K. platyops</i>	n = 1	47	—	—	—	—	11.9	—
<i>Au. garhi</i>	n = 1	27	—	11.6	—	—	—	—
<i>Au. africanus</i>	n	9	2	5	9	12	14	4
	mean	34.2	16.3	9.6	11.3	13.2	14.6	0.87
	range	30 - 37	14.5 - 18.0	8.8 - 10.0	10.8 - 12.4	12.4 - 14.4	13.3 - 16.4	0.84 - 0.89
	s.d.	1.9	—	0.48	0.50	0.50	0.95	0.02
early <i>Homo</i>	n	—	—	7	12	18	17	9
	mean	—	—	9.4	10.5	12.7	13.3	0.81
	range	—	—	8.3 - 11.5	9.7 - 12.1	11.4 - 14.5	12.5 - 16.0	0.77 - 0.88
	s.d.	—	—	1.35	0.68	0.84	0.96	0.03
<i>P. aethiopicus</i>	n	2	1	—	—	—	—	—
	mean	34	15	—	—	—	—	—
	range	31 - 37	—	—	—	—	—	—
	s.d.	—	—	—	—	—	—	—
<i>P. robustus</i>	n	8	5	13	18	19	16	7
	mean	36.8	13.8	8.5	12.7	14.1	15.2	0.89
	range	32 - 39	12.2 - 15.5	7.6 - 9.2	11.4 - 14.0	13.3 - 15.3	14.0 - 16.3	0.88 - 0.97
	s.d.	4.1	1.4	0.43	0.65	0.43	0.64	0.04
<i>P. boisei</i>	n	2	3	8	6	12	11	6
	mean	35.9	21	8.6	14.0	15.3	16.4	0.93
	range	33 - 39	20 - 22	6.5 - 10.8	13.5 - 14.6	13.3 - 16.4	15.6 - 19.0	0.89-0.94
	s.d.	—	1.0	1.18	0.54	1.02	1.43	0.02

Canine comparative measurements are taken from refs 13, 22, 23 and 26. Subnasal angle comparative measurements are taken from summary data compiled in ref. 4. Palate depth comparative measurements are from refs 3 and 27. All other data for the comparative taxa were compiled from refs 3, 23, 25, 28-33. Subnasal angle is measured as the subnasal clivus (nasospinale-prosthion) to the postcanine alveolar margin. Palate depth is measured at M¹/M² for all specimens, except for *P. aethiopicus* (EP 1500/01), which was measured at P³/P⁴ (but appears to retain a constant depth posteriorly²⁷).

Extended Data Table 3 | Enamel thickness measurements

MicroCT scan settings					
Fossil	Energy Settings	Field of View (mm)	x, pixel size (mm)	z slice thickness (mm)	
BRT-VP-3/1	180 kV, 0.250 mA	43.01	0.042	0.046	
BRT-VP-3/14	250 kV, 2.0 mA	131.69	0.129	0.129	
WYT-VP-2/10	250 kV, 1.0 mA	77.86	0.076	0.074	

BRT-VP-3/1 Two-dimensional enamel thickness						
	Cusp tip thickness (mm)	Maximum occlusal thickness (mm)		Maximum lateral thickness (mm)		Crown Area (mm ²)
M ¹	<i>Buccal</i> too worn	<i>Lingual</i> too worn	Buccal 1.11	Lingual 1.03	Buccal 1.2	<i>Lingual</i> 1.38
M ²	1.76	1.93	1.53	1.66	1.71	2.15

BRT-VP-3/1 Three-dimensional enamel thickness					
	Enamel Volume (mm ³)	Dentine area (mm ²)	Dentine volume (mm ³)	AET (average enamel thickness)	AETSTD (AET standardized)
M ² 3D	325.357	186.478	273.224	1.745	26.89

MicroCT scan settings, two-dimensional and three-dimensional enamel thickness measurements for BRT-VP-3/1.

Extended Data Table 4 | Three-dimensional and two-dimensional average and relative enamel thicknesses

	BRT-VP-3/1 M ²	<i>Homo sapiens</i>	<i>P. robustus</i>	<i>A. africanus</i>	<i>H. neanderthalensis</i>	<i>Pongo pygmaeus</i>	<i>Pan troglodytes</i>	<i>Gorilla gorilla</i>
3D Sample	1	39	9	9	29	12	26	9
Enamel volume (mm ³)	325.36	218.91	557.51	478.85	225.47	197.91	137	372.01
Dentine volume (mm ³)	186.48	226.83	507.6	598.9	345.65	336.05	166.05	1023.18
EDJ surface area (mm ²)	273.22	162.56	311.8	332.67	53.8	199.75	116.21	375.89
3D AET (mm)	1.74	1.43	1.83	1.48	1.08	1.01	0.75	0.98
3D RET	26.89	23.97	23.27	17.7	15.55	14.49	11.8	9.77
2D Sample	1	257	9	9	42	41	40	15
Enamel area (mm ²)	24.76	24.19	44.75	36.2	21.97	23.42	14.63	29.37
Dentine area (mm ²)	28.83	38.73	53.23	52.27	41.65	50.93	36.95	79.29
EDJ length (mm)	17.38	19.6	21.87	22.42	20.75	21.34	19.47	28.25
2D AET (mm)	1.42	1.22	2.03	1.63	1.06	1.1	0.75	1.04
2D RET	34.17	20.06	28.38	22.79	16.44	15.49	13.23	11.68

The values for the BRT-VP-3/1 M² are compared with extinct hominins and extant hominoids. Data for fossil hominins and hominoids taken from ref. 34. AET, average enamel thickness; EDJ, enamel–dentine junction; RET, relative enamel thickness.

Extended Data Table 5 | Mandibular dental dimensions (P_3 – M_3) and robusticity indices (at the M_1 and M_2 levels) of *Au. deyiremeda* and other Plio-Pleistocene hominins

	P_3		P_4		M_1		M_2		M_3		Robusticity Index at M_1	Robusticity Index at M_2
	MD	BL	MD	BL	MD	BL	MD	BL	MD	BL	-	-
	n	-	n	-	n	-	n	-	n	-	-	-
BRT-VP-3/14	n = 1	7.6	10.8	8.2	10.6	(12.6)	(12.6)	(14.8)	(13.8)	(15.8)	(13.4)	0.72
WYT-VP-2/10	n = 1	-	-	-	-	-	-	-	-	-	-	0.76
KNM-WT 8556	n = 1	9.9	12.2	11.4	12.8	13.6	13.1	-	(15.0)	(19.0)	-	(0.59)
LD 350-1	n = 1	-	-	8.7	10.5	(12.2)	-	13.2	12.5	13.0	12.2	(0.61)
<i>Au. anamensis</i>	n	6	5	6	7	8	8	7	7	7	2	-
	mean	9.9	10.9	9.0	10.6	12.6	11.7	14.3	13.5	15.0	13.2	0.53
	range	9.3 - 10.9	9.5 - 12.0	7.4 - 9.8	9.6 - 11.9	11.6 - 13.7	10.2 - 13.3	13.0 - 15.9	12.3 - 14.9	13.7 - 17.0	12.1 - 13.7	0.51 - 0.54
	s.d	0.60	1.00	0.95	0.84	0.87	1.10	0.92	0.88	1.17	0.66	-
<i>Au. afarensis</i>	n	29	29	26	22	31	24	33	31	26	23	19
	mean	9.5	10.7	9.8	11.0	13.1	12.6	14.3	13.5	15.2	13.4	0.58
	range	7.9 - 12.6	8.9 - 13.6	7.7 - 11.4	9.8 - 11.8	10.1 - 14.8	11.0 - 13.9	12.1 - 16.5	11.1 - 15.2	13.4 - 18.1	11.3 - 15.3	0.48 - 0.76
	s.d	1.00	1.10	1.00	0.84	0.90	0.80	1.20	1.00	1.30	1.00	0.07
<i>Au. africanus</i>	n	18	17	24	20	32	30	37	40	34	31	4
	mean	9.7	11.4	10.4	11.6	13.9	13.2	15.7	14.4	16.1	14.5	0.61
	range	8.4 - 11.2	9.2 - 13.9	9.3 - 12.3	10.3 - 13.4	12.4 - 15.8	10.8 - 15.1	14.0 - 17.8	12.7 - 16.8	12.9 - 18.5	12.1 - 16.8	0.56 - 0.65
	s.d	0.74	1.17	0.76	0.90	1.07	0.90	0.98	0.99	1.10	1.08	0.04
Early <i>Homo</i>	n	12	11	9	9	15	16	14	15	12	12	16
	mean	10.0	10.4	10.5	10.9	14.1	12.2	15.5	13.6	15.4	13.0	0.63
	range	8.9 - 11.1	8.0 - 12.3	9.1 - 11.8	8.1 - 12.7	13.0 - 15.6	10.6 - 14.1	13.9 - 18.3	11.7 - 15.4	14.0 - 17.5	11.7 - 14.7	0.51 - 0.74
	s.d	0.71	1.36	0.90	1.42	0.89	0.94	1.28	1.07	0.96	0.94	0.06
<i>P. robustus</i>	n	20	19	17	16	24	18	20	20	22	21	5
	mean	10.2	11.6	11.5	13.1	14.7	13.7	16.3	14.7	17.1	14.5	0.66
	range	9.0 - 11.4	9.0 - 13.7	10.6 - 12.6	11.5 - 14.8	13.2 - 16.5	11.8 - 15.0	14.8 - 17.9	12.8 - 16.3	15.1 - 20.5	12.6 - 17.0	0.55 - 0.75
	s.d	0.61	1.14	0.62	1.11	0.77	0.85	0.90	0.84	1.38	1.14	0.07
<i>P. boisei</i>	n	5	5	9	9	10	6	8	6	9	12	23
	mean	11.1	13.0	13.5	14.3	16.5	15.5	18.2	16.9	16.4	20.0	0.69
	range	8.9 - 13.0	11.4 - 13.7	10.1 - 15.6	12.3 - 16.5	15.4 - 18.6	14.4 - 17.6	16.4 - 20.0	15.8 - 18.6	14.7 - 19.2	17.6 - 22.4	0.57 - 0.80
	s.d	1.66	0.93	1.76	1.28	0.97	1.17	1.33	1.11	1.56	1.71	0.05
												0.63 - 0.86

The values reported for BRT-VP-3/14 are averages of the left and right sides. The data for LD 350-1 are taken from ref. 21 and corrected for interproximal wear. Summary statistics for the comparative taxa were taken from ref. 15 and the references listed in Extended Data Table 2. Values in parentheses are estimates or corrected for breakage and/or interproximal attrition.

New species from Ethiopia further expands Middle Pliocene hominin diversity

Yohannes Haile-Selassie^{1,2}, Luis Gibert³, Stephanie M. Melillo⁴, Timothy M. Ryan⁵, Mulugeta Alene⁶, Alan Deino⁷, Naomi E. Levin⁸, Gary Scott⁷ & Beverly Z. Saylor²

Middle Pliocene hominin species diversity has been a subject of debate over the past two decades, particularly after the naming of *Australopithecus bahrelghazali* and *Kenyanthropus platyops* in addition to the well-known species *Australopithecus afarensis*. Further analyses continue to support the proposal that several hominin species co-existed during this time period. Here we recognize a new hominin species (*Australopithecus deyiremeda* sp. nov.) from 3.3–3.5-million-year-old deposits in the Woranso-Mille study area, central Afar, Ethiopia. The new species from Woranso-Mille shows that there were at least two contemporaneous hominin species living in the Afar region of Ethiopia between 3.3 and 3.5 million years ago, and further confirms early hominin taxonomic diversity in eastern Africa during the Middle Pliocene epoch. The morphology of *Au. deyiremeda* also reinforces concerns related to dentognathic (that is, jaws and teeth) homoplasy in Plio-Pleistocene hominins, and shows that some dentognathic features traditionally associated with *Paranthropus* and *Homo* appeared in the fossil record earlier than previously thought.

Recent fossil hominin discoveries from eastern and central Africa have been described as indicating the existence of multiple locomotor adaptations¹ and taxonomic diversity^{2–5} during the Middle Pliocene. Notably, *Kenyanthropus platyops* from Kenya and *Australopithecus bahrelghazali* from Chad were argued to be morphologically distinct from the contemporaneous *Au. afarensis*^{2–5}. Concerns have been raised about the validity of these two taxa, either because identification was based largely on a single specimen or because the type specimen was heavily distorted⁶. Yet further studies addressing these concerns maintain the initial proposal of taxonomic diversity^{4,5}. More recently, we described a 3.4-million-year (Myr)-old partial foot (BRT-VP-2/73) from the Burtele area of Woranso-Mille, Ethiopia¹. The specimen is contemporaneous with *Au. afarensis*, but demonstrates the existence of a distinct mode of bipedal locomotion. The BRT-VP-2/73 foot combines a grasping hallux with lateral rays that exhibit adaptation to hyperdorsiflexion at the metatarsophalangeal joint, similar to the morphology described for *Ardipithecus ramidus*⁷. This specimen was not assigned to any taxon, pending the recovery of cranial and dentognathic fossils from the same geographic and stratigraphic vicinity. Continued fieldwork in the Burtele area has resulted in the recovery of new dentognathic fossil remains, including partial and complete mandibles, a partial maxilla with dentition, and few isolated teeth (Table 1). Unfortunately, none of them is clearly associated with BRT-VP-2/73 (see Supplementary Note 1 for details).

The new remains described here show clear similarities with *Australopithecus* and lack the diagnostic features of *Ardipithecus*, *Paranthropus*, *Kenyanthropus* and *Homo*. Various aspects of dentognathic morphology distinguish these specimens from the contemporaneous *Au. afarensis* and warrant their assignment to a new species (Fig. 1).

Order Primates Linnaeus, 1758

Suborder Anthropoidea Mivart, 1864

Superfamily Hominoidea Gray, 1825

Genus *Australopithecus* Dart, 1925

***Australopithecus deyiremeda* sp. nov.**

Etymology. From the local Afar language terms *deyi*, meaning close, and *remeda*, meaning relative; thus referring to the species being a close relative of all later hominins.

Holotype. The holotype is BRT-VP-3/1 (Fig. 1a–e), a left maxilla with upper second incisor–upper second molar (I^2 – M^2) found by M. Barao on 4 March 2011. The originals of the holotype and paratypes are housed at the Paleoanthropology Laboratory of the National Museum of Ethiopia, Addis Ababa.

Paratypes. The paratypes are BRT-VP-3/14 (Fig. 1f–h), a complete mandible corpus with the apical halves of left lower incisor–right lower incisor (LI_2 – RI_2) and crowns of right lower premolar–lower molar (RP_3 – M_3), lacking the ascending rami; and WYT-VP-2/10 (Fig. 1i–k), a right edentulous mandible. Both specimens were found by A. Asfaw. Detailed anatomical descriptions, additional images, and measurements are provided in Supplementary Note 2, Extended Data Figs 1 and 2, and Extended Data Tables 1, 2 and 5.

Referred specimen. BRT-VP-3/37 (Fig. 1l–p), a right maxillary fragment with P^4 . This specimen was found only 5 m east of BRT-VP-3/14 and both specimens may belong to the same individual. However, this cannot be determined with confidence owing to the fragmentary nature of BRT-VP-3/37.

Localities. The Burtele (BRT) and Waytaleya (WYT) collection areas are located south of the perennial Mille River and north of the Mille-Chifra road. All localities are situated within an area of 3 km² (Fig. 2). The coordinates of the holotype are 11° 27' 43.9'' N and 40° 31' 41.0'' E (GPS datum WGS84).

Provenience. The holotype and paratypes from BRT-VP-3 are surface finds, weathered out of a sandstone at the bottom of a 6-m thick section of sandstone and siltstone capped by a 2-m thick basaltic flow. A tuff radiometrically dated to 3.469 ± 0.008 (mean \pm s.d.) Myr is located approximately 1 m above the basaltic flow. The paratype from WYT-VP-2 was also a surface find, collected from the top of a thin sandstone layer about 15 m above the 3.469-Myr-old tuff. The combined evidence from radiometric, palaeomagnetic and

¹Cleveland Museum of Natural History, Cleveland, Ohio 44106, USA. ²Case Western Reserve University, Cleveland, Ohio 44106, USA. ³University of Barcelona, Martí Franquès s/n, 08028 Barcelona, Spain.

⁴Max Plank Institute for Evolutionary Anthropology, Deutscher Platz 6, D-04103 Leipzig, Germany. ⁵Pennsylvania State University, University Park, Pennsylvania 16802, USA. ⁶Addis Ababa University, PO Box 1176 Addis Ababa, Ethiopia. ⁷Berkeley Geochronology Center, 2455 Ridge Road, Berkeley, California 94709, USA. ⁸Johns Hopkins University, Baltimore, Maryland 21218, USA.

Table 1 | Fossil hominins from BRT and WYT localities

Specimen number	Preserved elements	Date discovered	Discoverer	Measurements
BRT-VP-1/1	Left M_3 crown fragment	20 February 2006	K. Helem	N/A
BRT-VP-1/2	Left P_4	20 February 2006	K. Helem	10.4 (13.0)
BRT-VP-1/13	Right C_1	20 February 2006	C. Mesfin	9.9, 9.3
BRT-VP-1/18	Left M^2 mesial half	20 February 2006	K. Kayranto	N/A, 13.8
BRT-VP-1/118	Left M_3 fragment	17 February 2012	G. Alemayehu	N/A
BRT-VP-1/120	Left P_4	17 February 2012	A. Elema	9.9, 10.1
BRT-VP-2/10	Frontal fragment	20 February 2006	Y. Haile-Selassie	N/A
BRT-VP-2/73	Right partial foot	15 February 2009	S. Melillo	See ref. ¹
BRT-VP-2/89	Right P^4	17 February 2012	H. Gebreyesus	9.6, 12.9
BRT-VP-2/104	Molar fragment	26 February 2013	M. Mekomeli	N/A
BRT-VP-3/1	Left maxilla (I^2-M^2)	4 March 2011	M. Barao	I^2 , 6.9, 6.6; LC^1 , (8.3), (9.4); LP^3 , 7.7, 12.0; LP^4 , 7.6, 11.8; LM^1 , 10.6, 11.0; LM^2 , 11.8, 13.2
BRT-VP-3/14	Mandible with dentition	4 March 2011	A. Asfaw	LI_1 , 5.4, 8.0; LI_2 , 6.5, 9.0; LC_1 , 8.0, 10.7; LP_3 , 7.6, 10.8; LP_4 , 8.2, 10.6; LM_1 , (12.6), (12.6); LM_2 , (14.8), (13.8); LM_3 , (15.8), (13.4)
BRT-VP-3/37	Right maxilla fragment (P^4)	1 March 2013	A. Asfaw	8.2 (12.1)
WYT-VP-2/10	Partial edentulous mandible	5 March 2011	A. Asfaw	N/A

Dental measurements are standard mesiodistal and buccolingual/labiolingual dimensions. Values in parentheses are corrected for breakage or interproximal attrition. All measurements are in millimetres and were taken on the original specimens by the authors. N/A, non-dental specimens and fragments with no measurement.

depositional rate analyses yields estimated minimum and maximum ages of approximately 3.3 and 3.5 Myr, respectively (Fig. 3, Extended Data Fig. 3 and Supplementary Note 3 for details).

Differential diagnosis. *Australopithecus deyiremeda* differs from *Ar. ramidus* by having relatively and absolutely thicker enamel on its molars; a P_4 with three-roots; a robust mandible. It differs from *Au. anamensis* in lacking an extremely receding mandibular symphysis; having a more robust mandibular corpus and a bicuspid P_3 . It is distinguished from *Au. afarensis* by its overall mandibular architecture, including the lack of a lateral corpus hollow; low and anteriorly

positioned origin of the ascending ramus with pronounced takeoff at the P_4/M_1 level. It also differs in features of the maxilla, including an anteriorly positioned zygomatic origin and in dental size and morphology, specifically an absolutely small upper postcanine dentition and an upper canine that is mesiodistally short and lacks lingual relief. *Australopithecus deyiremeda* differs from *Paranthropus* in lacking molarized premolars and reduced incisors. It is distinguished from *K. platyops* by its transversely convex subnasal region, with incisor alveoli considerably anterior to the bi-canine line; I^1 root larger and more rounded than I^2 root; more prognathic nasoalveolar clivus and

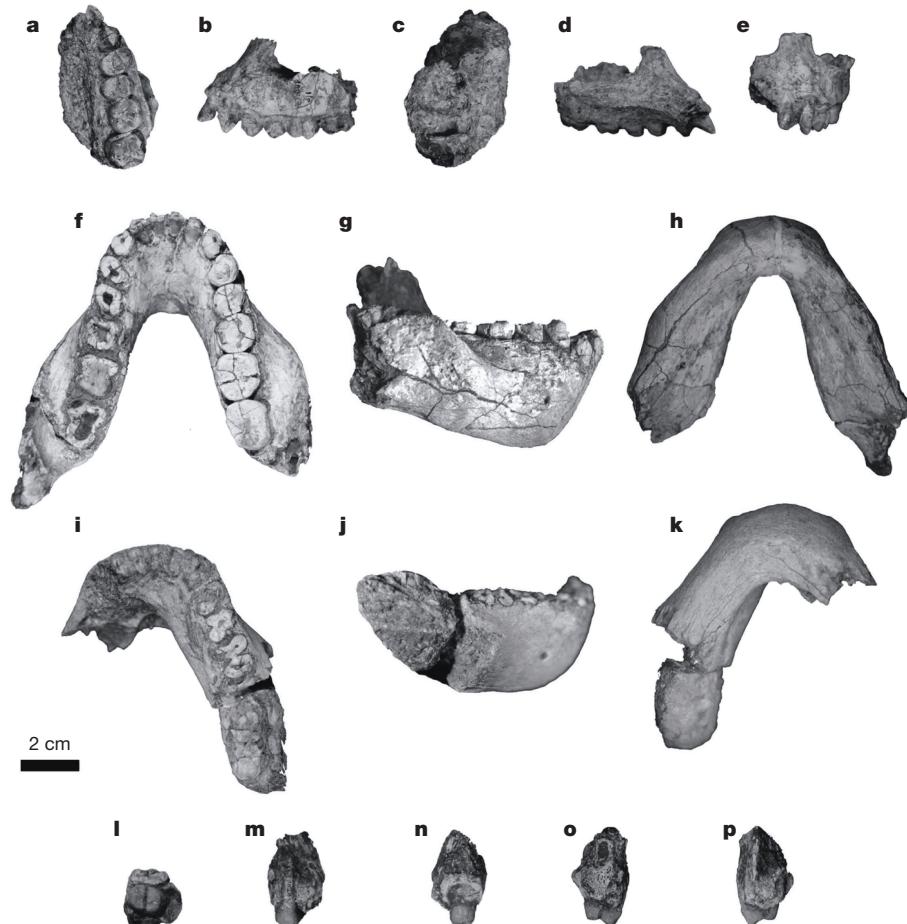


Figure 1 | Holotype BRT-VP-3/1. **a**, Occlusal view. **b**, Lateral view. **c**, Superior view. **d**, Medial view. **e**, Anterior view. Paratype BRT-VP-3/14. **f**, Occlusal view. **g**, Right lateral view. **h**, Basal view. Paratype WYT-VP-2/10. **i**, Occlusal view. **j**, Right lateral view. **k**, Basal view. Referred specimen BRT-VP-3/37. **l**, Occlusal view. **m**, Buccal view. **n**, Lingual view. **o**, Distal view. **p**, Mesial view.

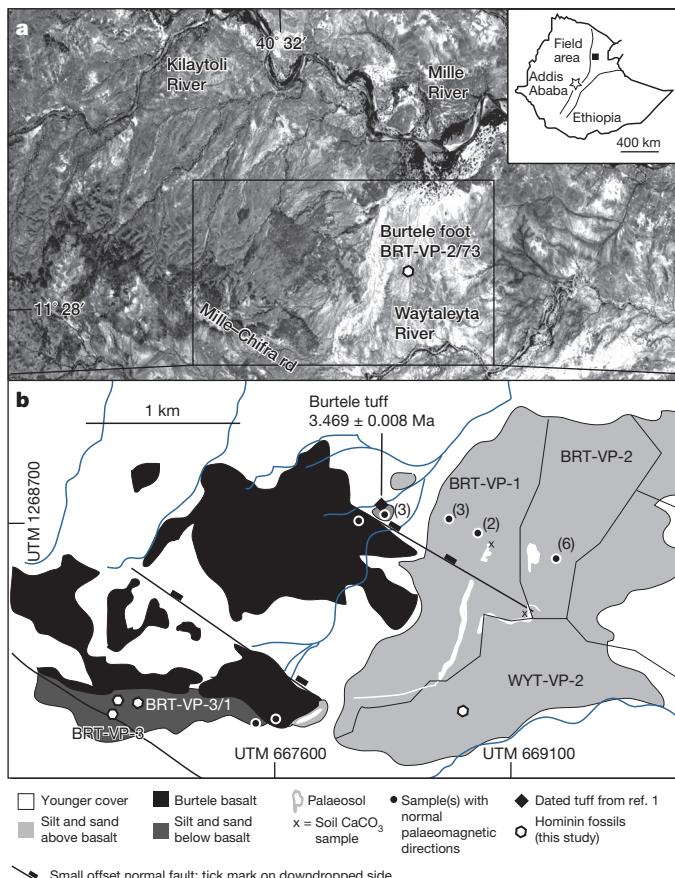


Figure 2 | Location map for the Burtele collection area and the BRT-VP-1, BRT-VP-2, BRT-VP-3 and WYT-VP-2 localities. **a**, Aerial photograph showing the location of the collection area in relation to nearby roads and rivers. **b**, Geological map of the study area showing the outcrop areas of sedimentary and volcanic units and the locations of fossil sites, palaeomagnetic samples (with number of samples indicated in parentheses), the dated Burtele tuff, and location of the holotype and paratypes of *Australopithecus deyiremeda*. The UTM (Universal Transverse Mercator) zone is 37P. The aerial photograph in **a** is reproduced with permission from the Ethiopian Mapping Agency.

I^2 root obliquely implanted, indicating some degree of procumbency; prominent upper canine and P^3 jugae; presence of incipient canine fossa and weakly expressed zygomaticoalveolar crest. *Australopithecus deyiremeda* can be distinguished from *Au. garhi* by its reduced subnasal prognathism and absolutely smaller canines and postcanine dentition. It differs from early *Homo* by its less incisiform canines, asymmetric P_3 , anterior margin of the mandibular ramus arising more anteriorly (opposite distal P_4), mandibular corpus height tapering posteriorly from P_3 to M_3 , and by its anterosuperiorly opening mental foramen.

Comparisons

The holotype maxilla (BRT-VP-3/1) has a more anteriorly positioned zygomatic root compared to *Au. anamensis* and most of the known *Au. afarensis* maxillae. However, it overlaps with most early hominins in the degree of subnasal prognathism (as determined by the subnasal angle) and palate depth^{3,4,8} (see Extended Data Table 2). Although the position of prosthion on BRT-VP-3/1 cannot be determined with certainty, the maximum estimated subnasal angle of this specimen is $\sim 39^\circ$, falling within the range of most early hominins but considerably less than the angle reported for KNM-WT 40000, the holotype of *K. platyops* (47°)⁴.

The incisive foramen in BRT-VP-3/1 is positioned at the distal bicanine line, whereas in *Au. afarensis* maxillae of comparable size (for

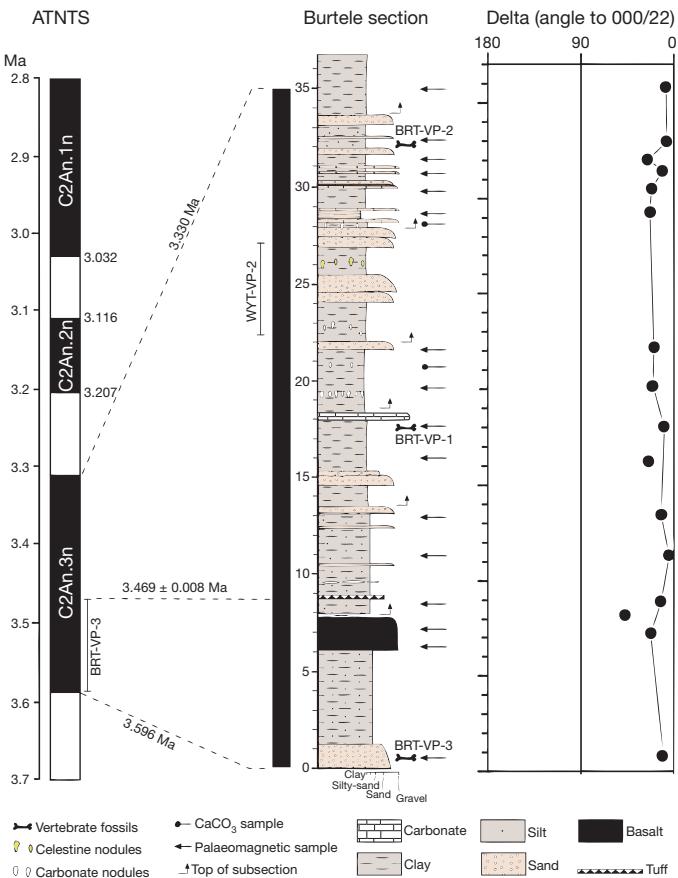


Figure 3 | Palaeomagnetic directions for 16 samples collected from the composite stratigraphic section at Burtele. Directions are shown as the angular distance (δ) to the expected normal direction ($000^\circ, 22^\circ$) for the given latitude. The stratigraphic position of the fossil localities, palaeomagnetic samples and carbonate samples are indicated. All palaeomagnetic samples exhibit normal remanence directions and fall within a single normal magnetozone. K-feldspar phenocrysts from a tuff ~ 1 m above the basalt have been dated by the $^{40}\text{Ar}/^{39}\text{Ar}$ method, yielding a mean age of 3.469 ± 0.008 Myr old¹, indicating that this 36-m section corresponds to the normal magnetozone C2An.3n (3.596–3.330 Ma). A minimum sedimentation accumulation rate of 18.6 cm per 1,000 years is calculated for 26 m of strata deposited in the time period <3.33 –3.47 Ma. Applying this rate, the fossil locality BRT-VP-3 situated 7 m stratigraphically below the tuff (excluding the 2 m basalt, see Supplementary Note 3) has an estimated age of 3.50–3.47 Myr.

example, A.L. 199-1), and early *Homo* specimens such as A.L. 666-1, the foramen is positioned posterior to this line, or at the P^3 level⁸. The palate anterior to the incisive foramen is also flexed inferiorly (see Extended Data Fig. 1i, k, l), a feature considered as derived although it is reported to be variable in *Au. afarensis*⁹. BRT-VP-3/1 shows very little sagittal overlap between the hard palate and the nasoalveolar clivus (3.5 mm) unlike most *Au. afarensis* maxillae (see Extended Data Fig. 4).

Further differences between *Au. deyiremeda* and *Au. afarensis* relate to the morphology of the upper dentition. The cheek teeth of *Au. afarensis* tend to exhibit roughly vertical buccal and lingual walls¹⁰, whereas the buccal and lingual walls of the *Au. deyiremeda* premolars and molars converge occlusally, making the occlusal surface narrower than the base of the crown. The P^3 and P^4 of BRT-VP-3/1 have three roots (Extended Data Fig. 1h, j). This is commonly seen in *Paranthropus* species, but highly variable in *Australopithecus*. *Kenyanthropus platyops* also has three-rooted upper premolars although P^3 – P^4 crown morphology is currently unknown for this species³.

The canine crown of BRT-VP-3/1 is absolutely smaller than almost all known *Australopithecus* upper canines, specifically in its

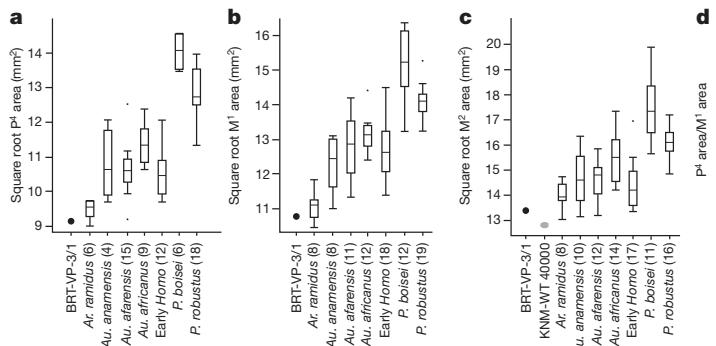


Figure 4 | Box-and-whisker diagrams comparing dental dimensions of the holotype BRT-VP-3/1 with other early hominins. Square roots of P⁴–M² areas (buccolingual \times mesiodistal dimensions) are plotted. The horizontal edges of each box indicate the first quartile (Q_1), the median (Q_2), and the third quartile (Q_3). The superiormost and inferiormost horizontal lines show the maximum and minimum values, respectively, excluding outliers, which are shown above or below these horizontal lines. Metric data for the comparative samples were compiled from refs 9 and 25, 28–33.

mesiodistal dimension, which falls close to the mean of *P. robustus* and *P. boisei* (Extended Data Table 2). It also appears to be long-rooted relative to its crown height, as in *Au. anamensis*^{11–13}. Furthermore, the canine of BRT-VP-3/1 does not have pronounced lingual relief and entirely lacks the prominent mesial and distal vertical lingual grooves commonly seen in *Au. afarensis* upper canines¹⁰. Its distal tubercle is also small and the mesial crown shoulder relatively short compared to the crown height and length of the mesial crest, thus it is less incisiform than *Au. afarensis* upper canines. It retains a more primitive, *Au. anamensis*-like condition in this regard.

The postcanine teeth of BRT-VP-3/1, particularly the P⁴ and M¹, are very small (Fig. 4; see Extended Data Table 2 for summary statistics). The M¹ area of BRT-VP-3/1 is the smallest of all known Pliocene hominins (less *Ar. ramidus*), whereas the P⁴ size is only approximated by A.L. 199-1, the smallest outlier in the *Au. afarensis* hypodigm. The M² also falls at the lower end of the *Au. afarensis* range and is slightly larger than that of KNM-WT 40000, the holotype of *K. platyops*. The P⁴ to M¹ area ratio of BRT-VP-3/1 falls within the range of most Pliocene hominin taxa, but outside the range of *Paranthropus*. In terms of enamel thickness, the two-dimensional linear enamel thickness for the M¹ of BRT-VP-3/1 is similar to that of other early hominins such as *Au. afarensis*, *P. robustus* and *Au. africanus*. However, the two-dimensional and three-dimensional enamel thickness values for the M² of this specimen are high, most similar to *P. robustus* (see Supplementary Note 4 and Extended Data Tables 3 and 4 for details).

Although *Au. deyiremeda* and *K. platyops* show some similarities in maxillary morphology and molar size, they can be distinguished by a suite of qualitative characters related to the morphology of the nasoalveolar clivus, shape of the dental arcade, and morphology of the anterior maxillary dentition. The nasoalveolar clivus in *K. platyops* is flat both sagittally and transversely³, whereas in *Au. deyiremeda* it is anteriorly convex in the transverse plane. In accordance with this difference, *Au. deyiremeda* and *K. platyops* differ in the shape of the anterior dental arcade. On the basis of the position of the preserved I¹ root, the transversely oblique orientation of the I² crown, and its placement relative to the canine and postcanine row, it is clear that the canine and incisors of BRT-VP-3/1 formed an arc with the incisor alveoli well anterior to the bi-canine line, as in *Au. afarensis* specimens such as A.L. 200-1a and A.L. 444-2 (ref. 9; see Extended Data Fig. 1f, g). In *K. platyops*, the incisors are arranged parallel to the bi-canine line as in *Paranthropus* and *Homo rudolfensis*³. *Kenyanthropus platyops* also has a derived low and curved zygomaticoalveolar crest³, whereas this feature is not sufficiently developed to form a visible crest in BRT-VP-3/1. Instead, the lateral alveolar wall blends smoothly with the base of the zygomatic bone. Furthermore, computed tomography scans show that the incisor roots of BRT-VP-3/1 are straight and inclined posterosuperiorly, indicating procumbent incisors (Extended Data Fig. 1i, k). By contrast, the incisor roots of *K. platyops* do not show any sign of procumbency³. On the basis of root cross-section approximately 2 mm below the buccal cervicoenamel line of I², the I¹ root of BRT-VP-3/1 is also rounded and more

robust than the mesiolaterally compressed I² root, unlike KNM-WT 40000 in which the I¹ and I² incisor roots are almost equal in size³ (see Extended Data Fig. 1m). In general, *K. platyops* seems to be more derived than *Au. deyiremeda* in most of its maxillary and dental morphology.

The overall morphology of the paratype mandibles (BRT-VP-3/14 and WYT-VP-2/10; Fig. 1 and Extended Data Fig. 2) is comparable to that of eastern African robust taxa, although their overall size is within the range of *Au. afarensis*. The anterior origin of the ascending ramus at the posterior P₄ level is comparable to *P. boisei* mandibles such as the Peninj mandible and KGA-10-525 (ref. 14), and *P. robustus* specimens such as SK-12. *Australopithecus afarensis* mandibles have the origin of the ascending ramus high and more posterior, mostly distal to the first molar^{9,15}. BRT-VP-3/14 and WYT-VP-2/10 have inflated mandibular corpora, especially at the M₂₋₃ level, and they lack the lateral corpus hollow ubiquitous in *Au. afarensis* mandibles⁹. Corpus dimensions of BRT-VP-3/14 and WYT-VP-2/10 are notable because they are broader for their size than specimens assigned to *Au. afarensis* or early *Homo*. In this respect, they are more similar to *Paranthropus* (Fig. 5 and Supplementary Note 5). Mandibular specimens from Lomekwi whose taxonomic affinities have not been established with certainty (KNM-WT 8556 and KNM-WT 16006)^{3,16}, are described as having robust corpora with an ascending ramus take-off at P₄ and M₁, respectively. However, the robusticity of the Lomekwi specimens is not as pronounced as it is in WYT-VP-2/10 and BRT-VP-3/14. Furthermore, the KNM-WT 8556 premolars are derived: the P₃ has a well-developed metaconid and the P₄ is large and molarized³. By contrast, the P₃ of BRT-VP-3/14 is almost unicuspied, with only a trace of the metaconid, and the P₄ is not molarized.

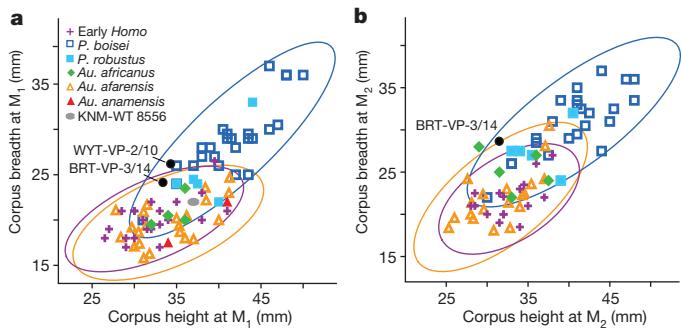


Figure 5 | Bivariate plots of mandibular corpus height and breadth among early hominins. a, Dimensions at the M₁ level. b, Dimensions at the M₂ level. The 95% confidence ellipses are shown for *Au. afarensis*, *P. boisei* and early *Homo* (see Supplementary Note 5 for details). The *Au. deyiremeda* corpora are relatively broad, comparable to *Paranthropus*. WYT-VP-2/10 falls outside the *Au. afarensis* ellipse at the M₁ level. BRT-VP-3/14 (represented by the mean of left and right sides) falls in an area of overlap at the M₁ level but at the edge of the *Au. afarensis* ellipse at the M₂ level. Mandibular measurements for the comparative taxa were taken from the references listed in Fig. 4 and ref. 14.

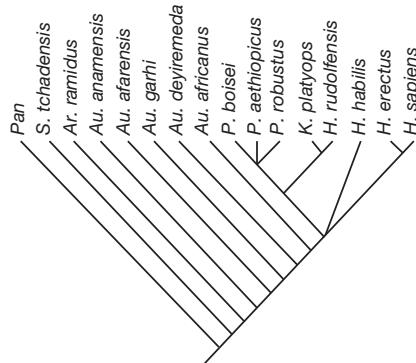


Figure 6 | Cladogram depicting the majority-rule consensus of 17 equally parsimonious trees that result from a phylogenetic analysis of the features preserved in the *Au. deyiremeda* hypodigm. Tree length = 63, consistency index = 0.63. The position of *Au. garhi* and topologies within the *Paranthropus* + *Kenyanthropus* + *Homo* clade are unstable when these features are considered. However, *Au. deyiremeda* is consistently placed as a sister taxon to a clade that includes *Au. africanus*, *Paranthropus*, and *Homo* (see Supplementary Notes 6–8 for further details).

The symphysis of BRT-VP-3/14 is thinner, superoinferiorly longer, and has a more receding external contour than that of WYT-VP-2/10 (Extended Data Fig. 2h, n). In the latter specimen, the superior half of the external contour is almost vertical but the inferior half is extremely receding similar to L.H.-4, the holotype of *Au. afarensis*^{9,17,18}. Further similarities between WYT-VP-2/10 and L.H.-4 include the prominence of the superior and inferior transverse tori (hence the deep genioglossal fossa) and the posterior extension of the inferior transverse torus exceeding that of the superior transverse torus. BRT-VP-3/14 has weak transverse tori and a very shallow genioglossal fossa (similar to *Au. afarensis* specimens such as A.L. 333w-12 and A.L. 315-22) and a pronounced postincisive planum (comparable to A.L. 400-1). These differences are within the range of variation seen in *Au. afarensis* mandibles⁹ and further show intra-taxon variability in mandibular symphyseal morphology. Although it was previously argued that two *Au. bahrelghazali* mandibles (KT 12 and KT 40) have symphyseal morphology that distinguishes them from *Au. afarensis* mandibles⁵, they appear to also fall within the range of *Au. afarensis* variation^{9,19,20}.

In terms of the dentition, BRT-VP-3/14 has small premolars compared to *Au. afarensis* (see Extended Data Table 5 and Supplementary Note 5 for details). The small mandibular premolars appear to conform well with the small upper premolars and M¹ of BRT-VP-3/1. Despite the robust appearance of the mandible, the BRT-VP-3/1 maxilla is only slightly smaller when occluded with the BRT-VP-3/14 mandible (see Extended Data Fig. 5).

The 2.8-Myr-old early *Homo* mandible (LD 350-1) recently reported from Ledi-Geraru²¹ is similar to *Au. deyiremeda* and some *Au. afarensis* specimens in its anterior corpus morphology. However, the derived features that are present in LD 350-1 and diagnostic of *Homo* are absent in the *Au. deyiremeda* mandibles (for example, posteriorly opening mental foramen, comparable anterior and posterior corpus height, M₃ mesiodistally shorter than M₂). Features of *Au. deyiremeda* that are derived relative to *Au. afarensis* and earlier hominins (for example, relatively robust mandibular corpus, anteriorly positioned root of the ascending ramus) are absent in LD 350-1 and other early *Homo* mandibles. Dentally, while the P₄ of BRT-VP-3/14 and LD 350-1 are comparable in size, the M₂ and M₃ of BRT-VP-3/14 are much larger (Extended Data Table 5).

Phylogenetic implications

The taxonomy and phylogenetic relationships among early hominins are becoming more complicated as new taxa are added to the Pliocene fossil record^{2,3,22,23} and the temporal range and systematics of early

Homo are reconsidered^{21,24}. A phylogenetic analysis of the preserved morphological features of *Au. deyiremeda* fails to produce a single most-parsimonious cladogram. However, the results are fairly consistent in placing this species as a sister taxon to a clade including *Au. africanus*, *Paranthropus* and *Homo* (Fig. 6 and Supplementary Notes 6–8). There are many alternative phylogenies consistent with this topology, but the fact that *Au. deyiremeda* is positioned outside the *Paranthropus* and *Homo* clade implies that some features associated with one or both of these taxa are homoplastic. Therefore, the addition of *Au. deyiremeda* reinforces questions about dentognathic homoplasy in later Pliocene and early Pleistocene hominins.

There is now incontrovertible evidence to show that multiple hominins existed contemporaneously in eastern Africa during the Middle Pliocene^{1,3}. Importantly, Woranso-Mille is geographically very close to Hadar (only 35 km north), where *Au. afarensis* is well documented. In combination, this suggests that multiple hominin species overlapped temporally and lived in close geographic proximity. What remains intriguing, and requires further investigation, is how these taxa are related to each other and to later hominins, and what environmental and ecological factors triggered such diversity. This has important implications for resource use and niche partitioning in Pliocene hominin palaeoecology, which can only be understood with the recovery and analysis of more hominin fossils and their associated fauna from Woranso-Mille and other contemporaneous sites.

Received 8 October 2014; accepted 9 April 2015.

- Haile-Selassie, Y. et al. A new hominin foot from Ethiopia shows multiple Pliocene bipedal adaptations. *Nature* **483**, 565–570 (2012).
- Brunet, M. et al. *Australopithecus bahrelghazali*, une nouvelle espèce d'Hominidé ancien de la région de Koro Toro (Tchad). *C. R. Acad. Sci.* **322**, 907–913 (1996).
- Leakey, M. G. et al. New hominin genus from eastern Africa shows diverse middle Pliocene lineages. *Nature* **410**, 433–440 (2001).
- Spoor, F., Leakey, M. G. & Leakey, L. N. Hominin diversity in the middle Pliocene of eastern Africa: the maxilla of KNM-WT 4000. *Phil. Trans. R. Soc. B* **365**, 3377–3388 (2010).
- Guy, F. et al. Symphyseal shape variation in extant and fossil hominoids, and the symphysis of *Australopithecus bahrelghazali*. *J. Hum. Evol.* **55**, 37–47 (2008).
- White, T. Early hominids—diversity or distortion? *Science* **299**, 1994–1997 (2003).
- Lovejoy, C. O., Latimer, B., Suwa, G., Asfaw, B. & White, T. D. Combining prehension and propulsion: the foot of *Ardipithecus ramidus*. *Science* **326**, 72e1–72e8 (2009).
- Kimbel, W. H., Johanson, D. C. & Rak, Y. Systematic assessment of a maxilla of *Homo* from Hadar, Ethiopia. *Am. J. Phys. Anthropol.* **103**, 235–262 (1997).
- Kimbel, W. H., Rak, Y. & Johanson, D. C. *The Skull of Australopithecus afarensis* (Oxford Univ. Press, 2004).
- Johanson, D. C., White, T. D. & Coppens, Y. Dental remains from the Hadar Formation, Ethiopia: 1974–1977 collections. *Am. J. Phys. Anthropol.* **57**, 545–603 (1982).
- Ward, C. V., Leakey, M. G. & Walker, A. Morphology of *Australopithecus anamensis* from Kanapoi and Allia Bay, Kenya. *J. Hum. Evol.* **41**, 255–368 (2001).
- Manthi, F., Plavcan, M. & Ward, C. New hominin fossils from Kanapoi, Kenya: mosaic evolution of canine evolution in early hominins. *So. Afr. J. Sci.* **108**, 1–9 (2012).
- Ward, C. V., Manthi, F. K. & Plavcan, J. M. New fossils of *Australopithecus anamensis* from Kanapoi, West Turkana, Kenya (2003–2008). *J. Hum. Evol.* **65**, 501–524 (2013).
- Suwa, G. et al. The first skull of *Australopithecus boisei*. *Nature* **389**, 489–492 (1997).
- Kimbel, W. H. & Delezenne, L. K. ‘Lucy’ redux: a review of research on *Australopithecus afarensis*. *Am. J. Phys. Anthropol.* **49**, 2–48 (2009).
- Brown, B., Brown, F. H. & Walker, A. New hominids from the Lake Turkana Basin, Kenya. *J. Hum. Evol.* **41**, 29–44 (2001).
- White, T. D. & Johanson, D. C. Pliocene hominid mandibles from the Hadar Formation, Ethiopia: 1974–1977 collections. *Am. J. Phys. Anthropol.* **57**, 501–544 (1982).
- Johanson, D. C., White, T. & Coppens, Y. *A New Species of the Genus Australopithecus (Primates: Hominidae) From the Pliocene of Eastern Africa* (Kirtlandia) (Cleveland Museum of Natural History, 1978).
- White, T. D., Suwa, G., Simpson, S. W. & Asfaw, B. Jaws and teeth of *Australopithecus afarensis* from Maka, Middle Awash, Ethiopia. *Am. J. Phys. Anthropol.* **111**, 45–68 (2000).
- Wood, B. & Richmond, B. Human evolution: taxonomy and paleobiology. *J. Anat.* **197**, 19–60 (2000).
- Villmoare, B. et al. Early *Homo* at 2.8 Ma from Ledi-Geraru, Afar, Ethiopia. *Science* **347**, 1352–1355 (2015).
- Asfaw, B. et al. *Australopithecus garhi*: a new species of early hominid from Ethiopia. *Science* **284**, 629–635 (1999).
- Leakey, M. G., Feibel, C. S., McDougall, I. & Walker, A. New four-million-year-old hominid species from Kanapoi and Allia Bay, Kenya. *Nature* **376**, 565–571 (1995).

24. Spoor, F. et al. Reconstructed *Homo habilis* type OH 7 suggests deep-rooted species diversity in early *Homo*. *Nature* **519**, 83–86 (2015).
25. Wood, B. *Koobi Fora Research Project Vol. 4* (Clarendon, 1991).
26. White, T. D., Johanson, D. C. & Kimbel, W. H. in *New Interpretations of Ape and Human Ancestry* (eds Ciochon, R. L. & Corruccini, R. S.) 721–780 (Plenum, 1983).
27. Harrison, T. in *Paleontology and Geology of Laetoli: Human Evolution in Context* (ed. Harrison, T.) Vol. 2 141–188 (Dordrecht: Springer, 2011).
28. White, T. D. Additional fossil hominids from Laetoli, Tanzania: 1976–1979 specimens. *Am. J. Phys. Anthropol.* **53**, 487–504 (1980).
29. White, T. D. New fossil hominids from Laetoli, Tanzania. *Am. J. Phys. Anthropol.* **46**, 197–229 (1977).
30. White, T. D. et al. Asa Issie, Aramis and the origin of *Australopithecus*. *Nature* **440**, 883–889 (2006).
31. White, T. D., Suwa, G., Simpson, S. W. & Asfaw, B. Jaws and teeth of *Australopithecus afarensis* from Maka. Middle Awash, Ethiopia. *Am. J. Phys. Anthropol.* **111**, 45–68 (2000).
32. Suwa, G. et al. Paleobiological implications of the *Ardipithecus ramidus* dentition. *Science* **326**, 94–99 (2009).
33. Leakey, M. G. et al. New fossils from Koobi Fora in northern Kenya confirm taxonomic diversity in early *Homo*. *Nature* **488**, 201–204 (2012).
34. Olejniczak, A. J. et al. Three-dimensional molar enamel distribution and thickness in *Australopithecus* and *Paranthropus*. *Biol. Lett.* **4**, 406–410 (2008).

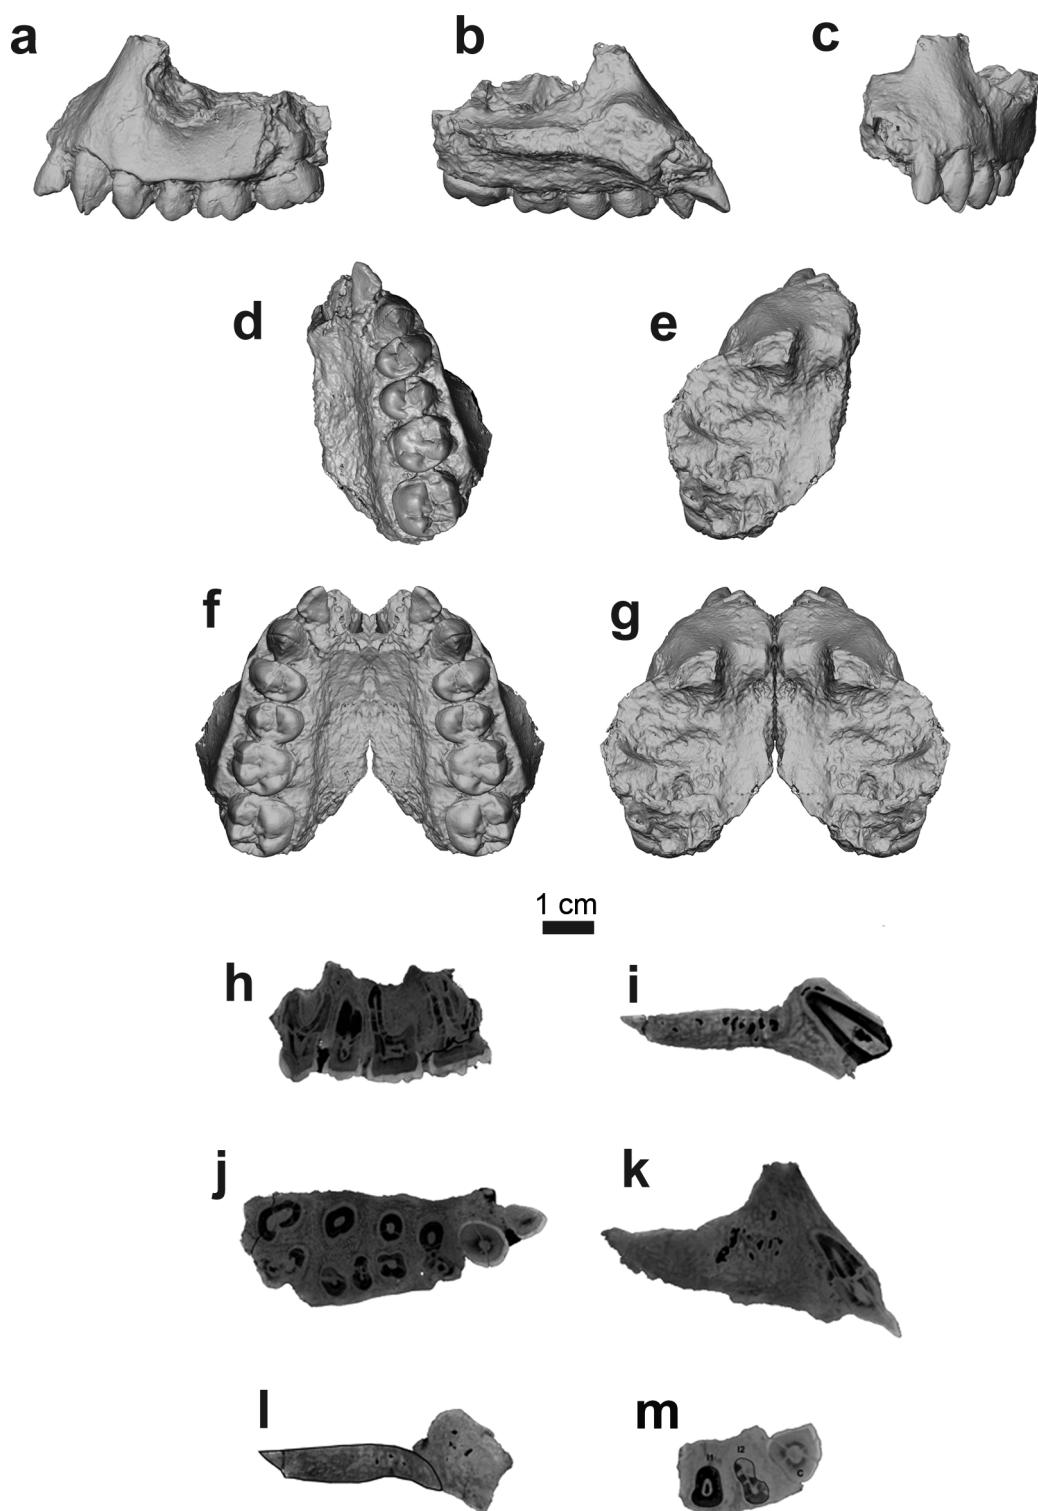
Supplementary Information is available in the online version of the paper.

Acknowledgements We thank the Authority for Research and Conservation of Cultural Heritage and the Afar Regional State of Ethiopia for permission to conduct field and

laboratory research, the Afar people of the Woranso–Mille area for their hospitality, and the project's fieldwork crew members for their incredible contributions in our fieldwork endeavours. We would also like to thank S. W. Simpson and W. H. Kimbel for their constructive comments and discussions throughout the preparation of this manuscript; G. Suwa, T. White and B. Asfaw for access to the original *Ardipithecus ramidus* material; W. H. Kimbel for access to the original *Australopithecus afarensis* material; F. Spoor for discussion of *Kenyanthropus platyops* morphology and Lomekwi mandibular metrics; D. F. Su for comments and assistance with the figures. This research was financially supported by grants from the LSB Leakey Foundation, the National Geographic Society, the Cleveland Museum of Natural History, and the National Science Foundation (BCS-0234320, BCS-0321893, BCS-0542037, BCS-1124705, BCS-1124713, BCS-1124716, BCS-1125157 and BCS-1125345).

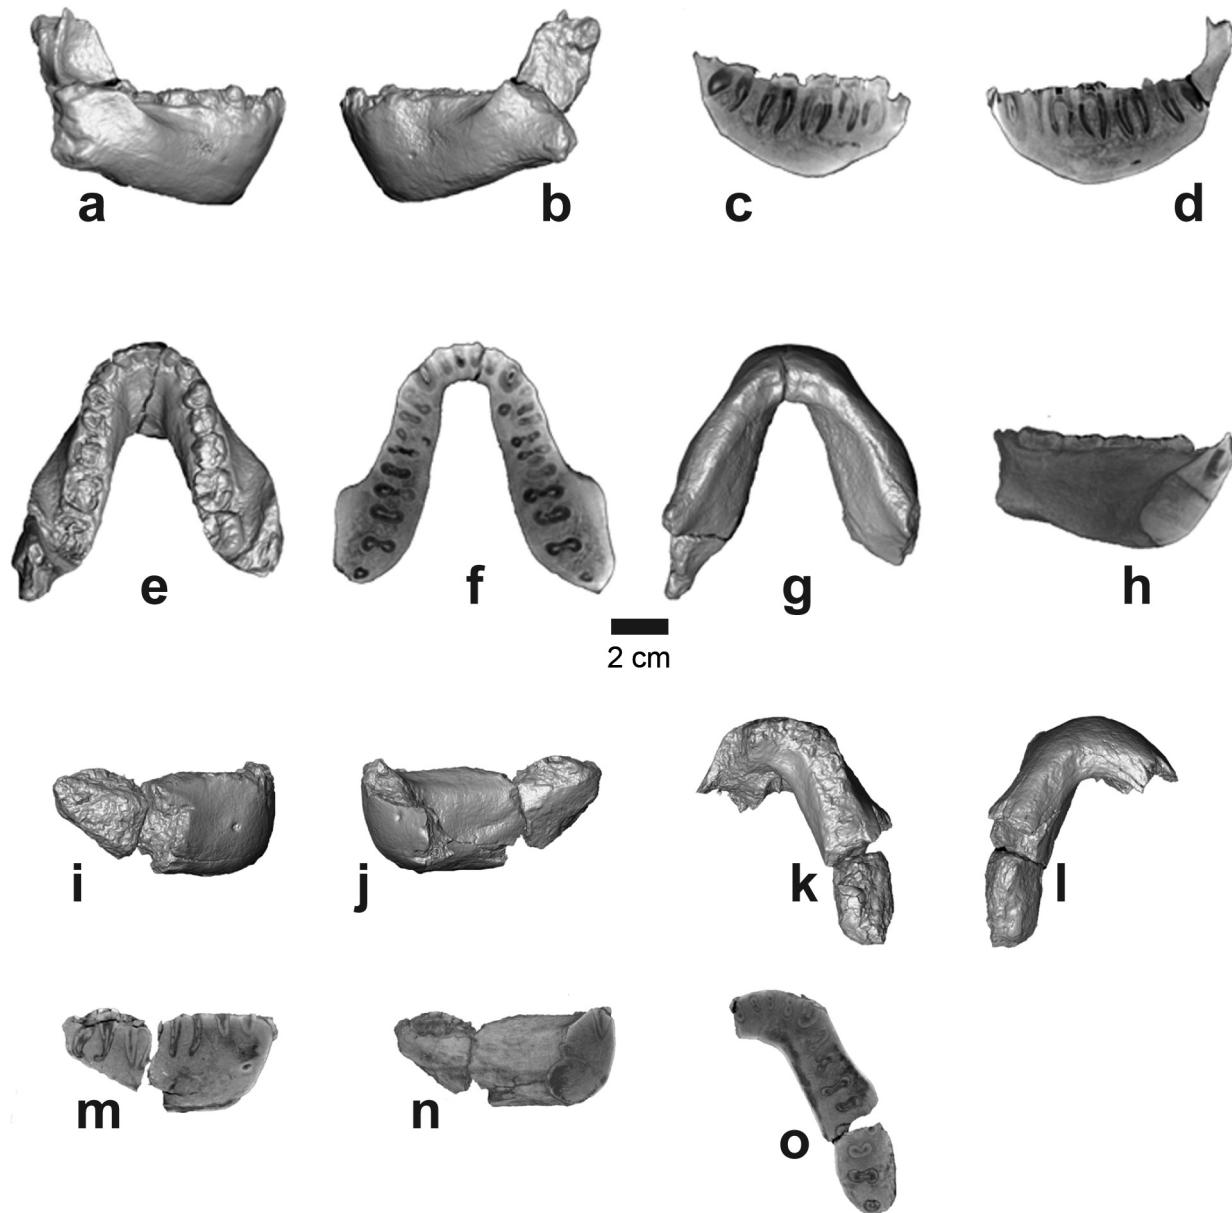
Author Contributions Y.H.-S. directed the field research. A.D., B.Z.S., M.A., L.G. and S.M.M. participated in the field research. T.M.R. conducted the computed tomography scanning and made the virtual reconstructions. A.D., B.Z.S., L.G. and M.A. studied the geological context, while L.G. and G.S. conducted the palaeomagnetic analysis. S.M.M. conducted the phylogenetic analysis. Y.H.-S. and S.M.M. made comparative observations and carried out analyses. Y.H.-S. took the lead in writing the paper with contributions from all coauthors.

Author Information The LSID urn:lsid:zoobank.org:pub:0C492889-01AC-4CDD-96FD-0E08A51F3CBB has been deposited in ZooBank. Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to Y.H.-S. (yhailese@cmnh.org).



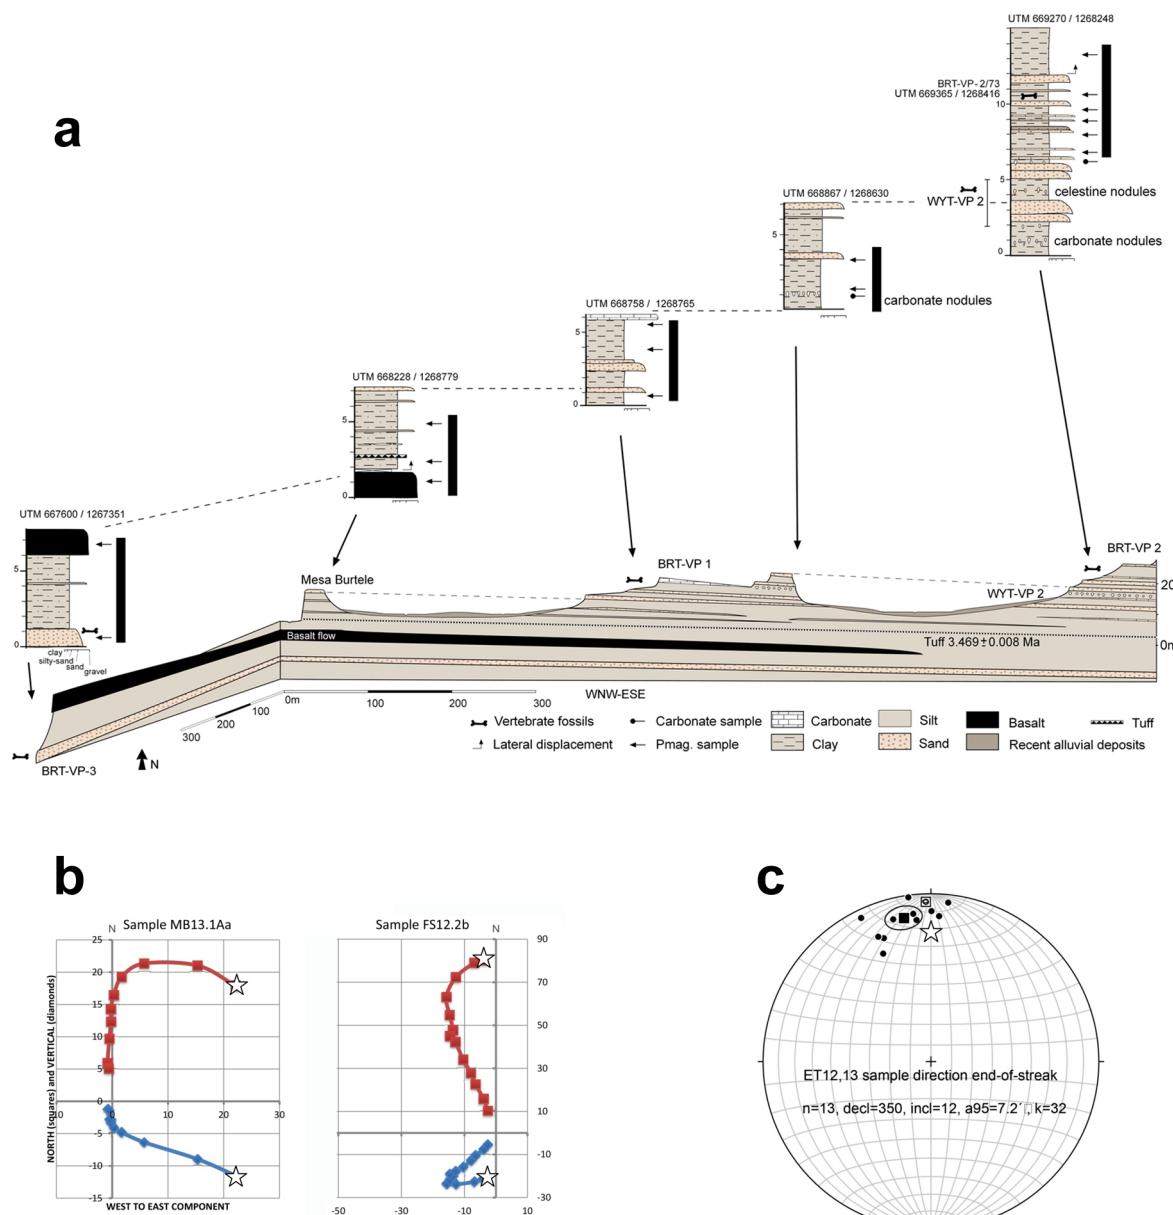
Extended Data Figure 1 | Computed-tomography-based visualization of BRT-VP-3/1. **a**, Lateral view. **b**, Medial view. **c**, Anterior view. **d**, Palatal view. **e**, Superior view. **f**, Palatal view mirror-imaged on midline. **g**, Superior view mirror-imaged on midline. Computed-tomography-based cross-sections of BRT-VP-3/1. **h**, Sagittal cross-section along the centre of the dental row

showing root morphology. **i**, Sagittal cross-section at I^1 . **j**, Transverse cross-section along the roots showing the number of roots of each tooth. **k**, Sagittal cross-section at I^2 . **l**, Midsagittal cross-section showing the shape of the palatine process and nasoalveolar clivus. **m**, Transverse cross-section across the incisors and the canine.



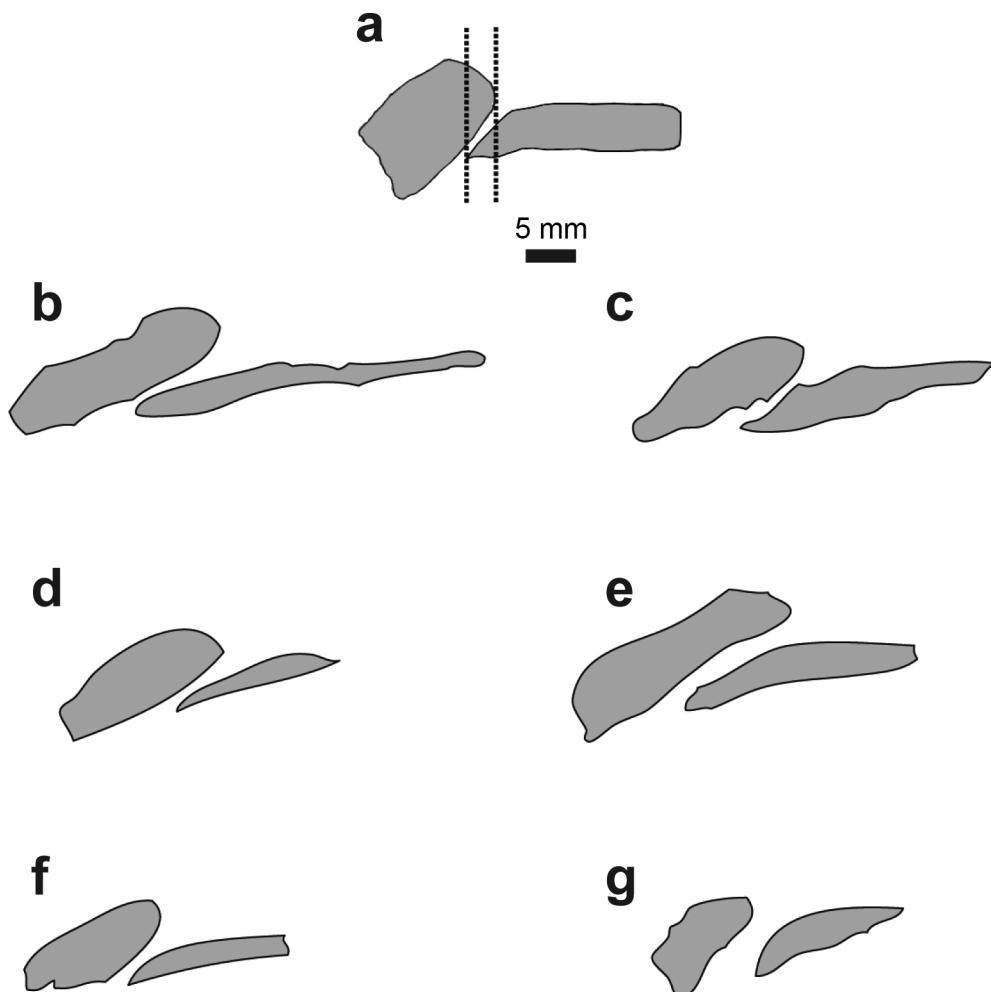
Extended Data Figure 2 | Computed-tomography-based visualization of BRT-VP-3/14. **a**, Right lateral view. **b**, Left lateral view. **c**, Sagittal cross-section along the centre of the right dental row showing root morphology. **d**, Sagittal cross-section along the centre of the left dental row showing root morphology. **e**, Occlusal view. **f**, Transverse cross-section along the roots showing the number of roots of each tooth. Note that the premolars have several roots.

g, Basal view. **h**, Symphyseal cross-section. Computed-tomography-based visualization of WYT-VP-2/10. **i**, Right lateral view. **j**, Medial view. **k**, Occlusal view. **l**, Basal view. **m**, Sagittal cross-section along the centre of the dental row showing root morphology. **n**, Symphyseal cross-section. **o**, Transverse cross-section along the roots showing the number of roots of each tooth.



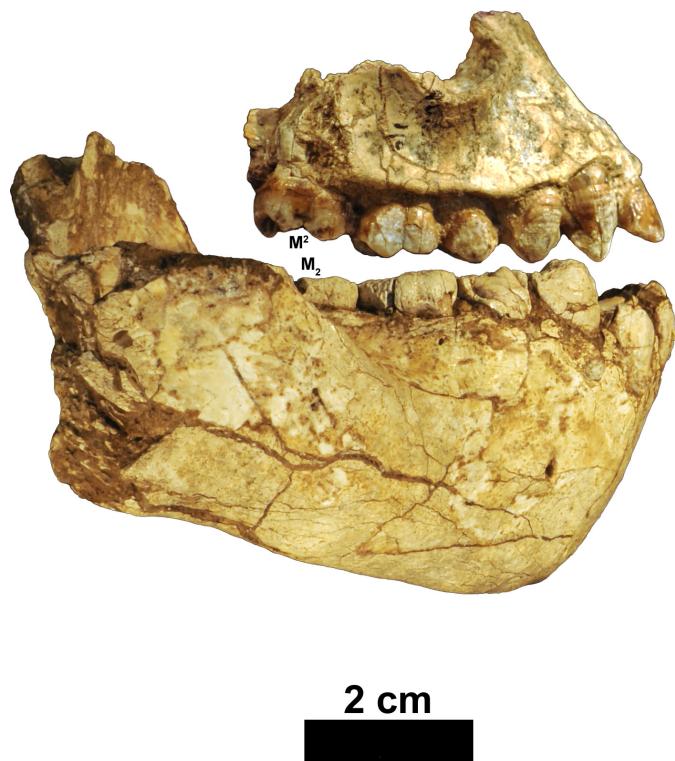
Extended Data Figure 3 | Magnetostratigraphy. **a**, Magnetostratigraphy for the Burtele area, measured in five subsections beginning with the oldest strata in the southwest (BRT-VP-3), to the youngest at BRT-VP-2 (see Fig. 2 for subsections locations). This stratigraphy consists predominantly of claystones/siltstones alternating with sandstones, and includes a 20-cm carbonate bed, a 2–6-m basalt flow, and a 5-cm tuff bed. The sandstones occupy slightly sinuous fluvial channels with a low width/depth ratio in a weakly confined setting. The sandstone beds are continuous and interstratified with pedogenically modified fine-grained overbank and ephemeral lake deposits. The proportion of sandstones increases upwards in the section, where it is accompanied by the occurrence of celestine nodules indicating increased aridity. Sixteen palaeomagnetic samples were obtained from fine-grained lithologies. All

samples demonstrate normal polarity (black bars), interpreted as belonging to a single polarity chron C2An.3n (3.596–3.330 Myr old) of the Astronomically Tuned Neogene Time Scale (ATNTS2004). **b**, Orthogonal demagnetization diagrams showing vector endpoints after alternating field and thermal demagnetization for samples MB13.1Aa and FS12.2b. Horizontal projections are shown as squares, vertical projection as diamonds, and NRM (4 mT) starting point as star. **c**, Palaeomagnetic directions for the Burtele stratigraphy from samples collected along different stratigraphic horizons. Stereographic projection referenced to magnetic North (declination 2° E), with solid symbols on lower hemisphere (plus inclination). Represented directions are the mean directions for each sample from three specimens. The start shows the location of the expected normal direction for this latitude (000/22).



Extended Data Figure 4 | Comparisons of midsagittal cross section of the hard palate and nasoalveolar clivus. a, BRT-VP-3/1. b, A.L. 444-2. c, A.L. 200-1. d, A.L. 486-1. e, A.L. 427-1. f, A.L. 199-1. g, A.L. 442-1. Note that the overlap

between the hard palate and the nasolalveolar clivus of BRT-VP-3/1 is small relative to most *Au. afarensis* specimens. Midsagittal cross-sections of *Au. afarensis* maxillae were modified from Fig. 5.22 in ref. 9.



Extended Data Figure 5 | BRT-VP-3/1 (reversed) and BRT-VP-3/14 shown in occlusion. The upper canine is aligned mesial to the P₃. Despite the apparently large size of the mandible (BRT-VP-3/14), the maxilla (BRT-VP-3/

1) is only slightly smaller. The position of the second molars is indicated to show that the canine-to-second-molar length is comparable in both specimens.

Extended Data Table 1 | Measurements of the paratype mandibles

Measurement	BRT-VP-3/14	WYT-VP-2/10
1 Symphyseal height	42.2	40.7
2 Maximum symphyseal depth	19.9	23.3
3 Horizontal position of mental foramen	P3/P4	P3
4 Mental foramen height from base	19.4	18
5 Mental foramen height from alveolar margin	~17.4	~16.6
6 Corpus height at P ₄	36.6, 35.8	35.7
7 Corpus width at P ₄	21.1, 21.8	23.8
8 Corpus height at P ₄ /M ₁	36.3, 35.9	34.8
9 Corpus width at P ₄ /M ₁	22.8, 22.7	27.11
10 Corpus height at M ₁	33.8, 33.0	34.3
11 Corpus width at M ₁	24.5, 23.6	26.2
12 Corpus height at M ₁ /M ₂	33.3, 31.9	-
13 Corpus width at M ₁ /M ₂	25.3, 28.0	-
14 Corpus height at M ₂	32.4, 30.5	-
15 Corpus width at M ₂	28.2, 29	-
16 Corpus height at M ₂ /M ₃	30.8, 28.5	-
17 Corpus width at M ₂ /M ₃	28.1, 29.5	-
18 P ₄ -P ₄ distance at mid-crown on lingual side	29.9	-
19 M ₁ -M ₁ distance at mid-crown on lingual side	~33.5	-

Corpus dimensions of BRT-VP-3/14 are reported for both right and left sides. Measurements 1–5 were taken following the methods used in ref. 25. Measurements 6–19 were taken following the methods used in ref. 17. All measurements are in millimetres.

Extended Data Table 2 | Statistical summary of measurements of early hominin maxillae and upper dentition

		Subnasal Angle	Palate Depth	C ¹ MD	$\sqrt{P^4}$ Area	$\sqrt{M^1}$ Area	$\sqrt{M^2}$ Area	$\sqrt{P^4}$ Area / $\sqrt{M^1}$ Area
BRT-VP-3/1		~39	~13.4	8.3	9.1	10.8	12.5	0.84
<i>Au. anamensis</i>	n	1	—	2	4	8	10	4
	mean	27	—	10.8	10.6	12.5	13.7	0.90
	range	—	—	9.9 - 11.7	9.7 - 12.1	11.0 - 13.1	12.3 - 15.5	0.79 - 0.92
	s.d.	—	—	—	1.0	0.74	1.02	0.06
<i>Au. afarensis</i>	n	6	6	10	15	11	12	11
	mean	34.6	11.3	9.9	10.6	12.9	13.9	0.81
	range	29 - 39	8.5 - 14	8.9 - 11.6	9.2 - 12.5	11.4 - 14.2	12.3 - 14.9	0.74 - 0.97
	s.d.	3.5	1.8	0.74	0.72	0.88	0.73	0.07
<i>K. platyops</i>	n = 1	47	—	—	—	—	11.9	—
<i>Au. garhi</i>	n = 1	27	—	11.6	—	—	—	—
<i>Au. africanus</i>	n	9	2	5	9	12	14	4
	mean	34.2	16.3	9.6	11.3	13.2	14.6	0.87
	range	30 - 37	14.5 - 18.0	8.8 - 10.0	10.8 - 12.4	12.4 - 14.4	13.3 - 16.4	0.84 - 0.89
	s.d.	1.9	—	0.48	0.50	0.50	0.95	0.02
early <i>Homo</i>	n	—	—	7	12	18	17	9
	mean	—	—	9.4	10.5	12.7	13.3	0.81
	range	—	—	8.3 - 11.5	9.7 - 12.1	11.4 - 14.5	12.5 - 16.0	0.77 - 0.88
	s.d.	—	—	1.35	0.68	0.84	0.96	0.03
<i>P. aethiopicus</i>	n	2	1	—	—	—	—	—
	mean	34	15	—	—	—	—	—
	range	31 - 37	—	—	—	—	—	—
	s.d.	—	—	—	—	—	—	—
<i>P. robustus</i>	n	8	5	13	18	19	16	7
	mean	36.8	13.8	8.5	12.7	14.1	15.2	0.89
	range	32 - 39	12.2 - 15.5	7.6 - 9.2	11.4 - 14.0	13.3 - 15.3	14.0 - 16.3	0.88 - 0.97
	s.d.	4.1	1.4	0.43	0.65	0.43	0.64	0.04
<i>P. boisei</i>	n	2	3	8	6	12	11	6
	mean	35.9	21	8.6	14.0	15.3	16.4	0.93
	range	33 - 39	20 - 22	6.5 - 10.8	13.5 - 14.6	13.3 - 16.4	15.6 - 19.0	0.89-0.94
	s.d.	—	1.0	1.18	0.54	1.02	1.43	0.02

Canine comparative measurements are taken from refs 13, 22, 23 and 26. Subnasal angle comparative measurements are taken from summary data compiled in ref. 4. Palate depth comparative measurements are from refs 3 and 27. All other data for the comparative taxa were compiled from refs 3, 23, 25, 28-33. Subnasal angle is measured as the subnasal clivus (nasospinale-prosthion) to the postcanine alveolar margin. Palate depth is measured at M¹/M² for all specimens, except for *P. aethiopicus* (EP 1500/01), which was measured at P³/P⁴ (but appears to retain a constant depth posteriorly²⁷).

Extended Data Table 3 | Enamel thickness measurements

MicroCT scan settings					
Fossil	Energy Settings	Field of View (mm)	x, pixel size (mm)	z slice thickness (mm)	
BRT-VP-3/1	180 kV, 0.250 mA	43.01	0.042	0.046	
BRT-VP-3/14	250 kV, 2.0 mA	131.69	0.129	0.129	
WYT-VP-2/10	250 kV, 1.0 mA	77.86	0.076	0.074	

BRT-VP-3/1 Two-dimensional enamel thickness						
	Cusp tip thickness (mm)	Maximum occlusal thickness (mm)		Maximum lateral thickness (mm)		Crown Area (mm ²)
M ¹	<i>Buccal</i> too worn	<i>Lingual</i> too worn	Buccal 1.11	Lingual 1.03	Buccal 1.2	<i>Lingual</i> 1.38
M ²	1.76	1.93	1.53	1.66	1.71	2.15

BRT-VP-3/1 Three-dimensional enamel thickness					
	Enamel Volume (mm ³)	Dentine area (mm ²)	Dentine volume (mm ³)	AET (average enamel thickness)	AETSTD (AET standardized)
M ² 3D	325.357	186.478	273.224	1.745	26.89

MicroCT scan settings, two-dimensional and three-dimensional enamel thickness measurements for BRT-VP-3/1.

Extended Data Table 4 | Three-dimensional and two-dimensional average and relative enamel thicknesses

	BRT-VP-3/1 M ²	<i>Homo sapiens</i>	<i>P. robustus</i>	<i>A. africanus</i>	<i>H. neanderthalensis</i>	<i>Pongo pygmaeus</i>	<i>Pan troglodytes</i>	<i>Gorilla gorilla</i>
3D Sample	1	39	9	9	29	12	26	9
Enamel volume (mm ³)	325.36	218.91	557.51	478.85	225.47	197.91	137	372.01
Dentine volume (mm ³)	186.48	226.83	507.6	598.9	345.65	336.05	166.05	1023.18
EDJ surface area (mm ²)	273.22	162.56	311.8	332.67	53.8	199.75	116.21	375.89
3D AET (mm)	1.74	1.43	1.83	1.48	1.08	1.01	0.75	0.98
3D RET	26.89	23.97	23.27	17.7	15.55	14.49	11.8	9.77
2D Sample	1	257	9	9	42	41	40	15
Enamel area (mm ²)	24.76	24.19	44.75	36.2	21.97	23.42	14.63	29.37
Dentine area (mm ²)	28.83	38.73	53.23	52.27	41.65	50.93	36.95	79.29
EDJ length (mm)	17.38	19.6	21.87	22.42	20.75	21.34	19.47	28.25
2D AET (mm)	1.42	1.22	2.03	1.63	1.06	1.1	0.75	1.04
2D RET	34.17	20.06	28.38	22.79	16.44	15.49	13.23	11.68

The values for the BRT-VP-3/1 M² are compared with extinct hominins and extant hominoids. Data for fossil hominins and hominoids taken from ref. 34. AET, average enamel thickness; EDJ, enamel–dentine junction; RET, relative enamel thickness.

Extended Data Table 5 | Mandibular dental dimensions (P_3 – M_3) and robusticity indices (at the M_1 and M_2 levels) of *Au. deyiremeda* and other Plio-Pleistocene hominins

	P_3		P_4		M_1		M_2		M_3		Robusticity Index at M_1	Robusticity Index at M_2
	MD	BL	MD	BL	MD	BL	MD	BL	MD	BL	-	-
	n	-	n	-	n	-	n	-	n	-	-	-
BRT-VP-3/14	n = 1	7.6	10.8	8.2	10.6	(12.6)	(12.6)	(14.8)	(13.8)	(15.8)	(13.4)	0.72
WYT-VP-2/10	n = 1	-	-	-	-	-	-	-	-	-	-	0.76
KNM-WT 8556	n = 1	9.9	12.2	11.4	12.8	13.6	13.1	-	(15.0)	(19.0)	-	(0.59)
LD 350-1	n = 1	-	-	8.7	10.5	(12.2)	-	13.2	12.5	13.0	12.2	(0.61)
<i>Au. anamensis</i>	n	6	5	6	7	8	8	7	7	7	2	-
	mean	9.9	10.9	9.0	10.6	12.6	11.7	14.3	13.5	15.0	13.2	0.53
	range	9.3 - 10.9	9.5 - 12.0	7.4 - 9.8	9.6 - 11.9	11.6 - 13.7	10.2 - 13.3	13.0 - 15.9	12.3 - 14.9	13.7 - 17.0	12.1 - 13.7	0.51 - 0.54
	s.d	0.60	1.00	0.95	0.84	0.87	1.10	0.92	0.88	1.17	0.66	-
<i>Au. afarensis</i>	n	29	29	26	22	31	24	33	31	26	23	19
	mean	9.5	10.7	9.8	11.0	13.1	12.6	14.3	13.5	15.2	13.4	0.58
	range	7.9 - 12.6	8.9 - 13.6	7.7 - 11.4	9.8 - 11.8	10.1 - 14.8	11.0 - 13.9	12.1 - 16.5	11.1 - 15.2	13.4 - 18.1	11.3 - 15.3	0.48 - 0.76
	s.d	1.00	1.10	1.00	0.84	0.90	0.80	1.20	1.00	1.30	1.00	0.07
<i>Au. africanus</i>	n	18	17	24	20	32	30	37	40	34	31	4
	mean	9.7	11.4	10.4	11.6	13.9	13.2	15.7	14.4	16.1	14.5	0.61
	range	8.4 - 11.2	9.2 - 13.9	9.3 - 12.3	10.3 - 13.4	12.4 - 15.8	10.8 - 15.1	14.0 - 17.8	12.7 - 16.8	12.9 - 18.5	12.1 - 16.8	0.56 - 0.65
	s.d	0.74	1.17	0.76	0.90	1.07	0.90	0.98	0.99	1.10	1.08	0.04
Early <i>Homo</i>	n	12	11	9	9	15	16	14	15	12	12	16
	mean	10.0	10.4	10.5	10.9	14.1	12.2	15.5	13.6	15.4	13.0	0.63
	range	8.9 - 11.1	8.0 - 12.3	9.1 - 11.8	8.1 - 12.7	13.0 - 15.6	10.6 - 14.1	13.9 - 18.3	11.7 - 15.4	14.0 - 17.5	11.7 - 14.7	0.51 - 0.74
	s.d	0.71	1.36	0.90	1.42	0.89	0.94	1.28	1.07	0.96	0.94	0.06
<i>P. robustus</i>	n	20	19	17	16	24	18	20	20	22	21	5
	mean	10.2	11.6	11.5	13.1	14.7	13.7	16.3	14.7	17.1	14.5	0.66
	range	9.0 - 11.4	9.0 - 13.7	10.6 - 12.6	11.5 - 14.8	13.2 - 16.5	11.8 - 15.0	14.8 - 17.9	12.8 - 16.3	15.1 - 20.5	12.6 - 17.0	0.55 - 0.75
	s.d	0.61	1.14	0.62	1.11	0.77	0.85	0.90	0.84	1.38	1.14	0.07
<i>P. boisei</i>	n	5	5	9	9	10	6	8	6	9	12	23
	mean	11.1	13.0	13.5	14.3	16.5	15.5	18.2	16.9	16.4	20.0	0.69
	range	8.9 - 13.0	11.4 - 13.7	10.1 - 15.6	12.3 - 16.5	15.4 - 18.6	14.4 - 17.6	16.4 - 20.0	15.8 - 18.6	14.7 - 19.2	17.6 - 22.4	0.57 - 0.80
	s.d	1.66	0.93	1.76	1.28	0.97	1.17	1.33	1.11	1.56	1.71	0.05
												0.63 - 0.86

The values reported for BRT-VP-3/14 are averages of the left and right sides. The data for LD 350-1 are taken from ref. 21 and corrected for interproximal wear. Summary statistics for the comparative taxa were taken from ref. 15 and the references listed in Extended Data Table 2. Values in parentheses are estimates or corrected for breakage and/or interproximal attrition.

Whole-genome characterization of chemoresistant ovarian cancer

Ann-Marie Patch^{1,2*}, Elizabeth L. Christie^{3*}, Dariush Etemadmoghadam^{3,4,5*}, Dale W. Garsed^{3*}, Joshy George⁶, Sian Fereday³, Katia Nones^{1,2}, Prue Cowin³, Kathryn Alsop³, Peter J. Bailey^{1,7}, Karin S. Kassahn^{1,8}, Felicity Newell¹, Michael C. J. Quinn^{1,2}, Stephen Kazakoff^{1,2}, Kelly Quek¹, Charlotte Wilhelm-Benartzl⁹, Ed Curry⁹, Huei San Leong³, The Australian Ovarian Cancer Study Group†, Anne Hamilton^{3,10,11}, Linda Mileshkin^{3,5}, George Au-Yeung³, Catherine Kennedy¹², Jillian Hung¹², Yoke-Eng Chiew¹², Paul Harnett¹³, Michael Friedlander¹⁴, Michael Quinn¹¹, Jan Pyman¹¹, Stephen Cordner¹⁵, Patricia O'Brien¹⁵, Jodie Leditschke¹⁵, Greg Young¹⁵, Kate Strachan¹⁵, Paul Waring⁴, Walid Azar³, Chris Mitchell³, Nadia Traficante³, Joy Hendley³, Heather Thorne³, Mark Shackleton^{3,5}, David K. Miller¹, Gisela Mir Arnau³, Richard W. Tothill^{3,5}, Timothy P. Holloway³, Timothy Semple³, Ivon Harliwong¹, Craig Nourse¹, Ehsan Nourbakhsh¹, Suzanne Manning¹, Senel Idrisoglu¹, Timothy J. C. Bruxner¹, Angelika N. Christ¹, Barsha Poudel¹, Oliver Holmes^{1,2}, Matthew Anderson¹, Conrad Leonard^{1,2}, Andrew Lonie¹⁶, Nathan Hall¹⁷, Scott Wood^{1,2}, Darrin F. Taylor¹, Qinying Xu^{1,2}, J. Lynn Fink¹, Nick Waddell¹, Ronny Drapkin¹⁸, Euan Stronach⁹, Hani Gabra⁹, Robert Brown⁹, Andrea Jewell¹⁹, Shivashankar H. Nagaraj¹, Emma Markham¹, Peter J. Wilson¹, Jason Ellul³, Orla McNally¹², Maria A. Doyle³, Ravikiran Vedururu³, Collin Stewart²⁰, Ernst Lengyel¹⁹, John V. Pearson^{1,2}, Nicola Waddell^{1,2}, Anna deFazio^{12§}, Sean M. Grimmond^{1,7§} & David D. L. Bowtell^{3,4,5,9,21§}

Patients with high-grade serous ovarian cancer (HGSC) have experienced little improvement in overall survival, and standard treatment has not advanced beyond platinum-based combination chemotherapy, during the past 30 years. To understand the drivers of clinical phenotypes better, here we use whole-genome sequencing of tumour and germline DNA samples from 92 patients with primary refractory, resistant, sensitive and matched acquired resistant disease. We show that gene breakage commonly inactivates the tumour suppressors *RBI*, *NFI*, *RAD51B* and *PTEN* in HGSC, and contributes to acquired chemotherapy resistance. *CCNE1* amplification was common in primary resistant and refractory disease. We observed several molecular events associated with acquired resistance, including multiple independent reversions of germline *BRCA1* or *BRCA2* mutations in individual patients, loss of *BRCA1* promoter methylation, an alteration in molecular subtype, and recurrent promoter fusion associated with overexpression of the drug efflux pump *MDR1*.

Recurrent, chemotherapy-resistant HGSC¹ accounts for most epithelial ovarian cancer deaths. Primary disease is characterized by ubiquitous *TP53* mutation^{2,3}, extensive copy number change^{3–5}, and in about half of all HGSCs, there is evidence of mutational³ and functional⁶ inactivation of homologous recombination repair. Extensive intratumoural heterogeneity in primary HGSCs has been recently documented, as it has in other solid cancers^{7–9}. However, relatively little is known of the genomic evolution of HGSC under the selective pressure of chemotherapy.

As HGSC is driven by genomic copy number change rather than recurrent point mutation⁵, we undertook the first whole-genome sequence (WGS) analysis of a large HGSC cohort to examine structural variation at high resolution. We performed transcriptome, methylation, and microRNA (miRNA) expression analyses to support the WGS data, and compared resistant, refractory and sensitive

primary disease. We also evaluated paired primary sensitive and relapse-resistant samples, and end-stage HGSC obtained from patients who underwent rapid autopsy for research.

Driver mutation by gene breakage

We performed whole genome sequencing (40× and 52× average coverage, germline and tumour, respectively; Supplementary Table 1), complemented by transcriptome sequencing (average 210 million reads per sample; Supplementary Table 2), and methylation, copy number and miRNA array analysis on a total of 114 tumour samples from 92 patients (Methods and Extended Data Fig. 1). Confirming previous reports^{2,3}, *TP53* mutations were prevalent (Fig. 1 and Supplementary Table 3). Inactivating germline or somatic mutations in genes associated with homologous recombination repair, or *BRCA1* methylation, were detected collectively in half of the primary tumours (Fig. 1 and

¹Queensland Centre for Medical Genomics, Institute for Molecular Bioscience, The University of Queensland, Brisbane, Queensland 4067, Australia. ²QIMR Berghofer Medical Research Institute, Brisbane, Queensland 4006, Australia. ³Peter MacCallum Cancer Centre, East Melbourne, Victoria 3002, Australia. ⁴Department of Pathology, University of Melbourne, Parkville, Victoria 3052, Australia. ⁵Sir Peter MacCallum Cancer Centre Department of Oncology, University of Melbourne, Parkville, Victoria 3052, Australia. ⁶The Jackson Laboratory for Genomic Medicine, Farmington, Connecticut 06030, USA. ⁷Wolfson Wohl Cancer Research Centre, Institute of Cancer Sciences, University of Glasgow, Glasgow G61 1QH, UK. ⁸Technology Advancement Unit, Genetics and Molecular Pathology, SA Pathology, Adelaide, South Australia 5000, Australia. ⁹Ovarian Cancer Action Research Centre, Department of Surgery and Cancer, Imperial College London, London W12 0HS, UK. ¹⁰Department of Medicine, University of Melbourne, Parkville, Victoria 3052, Australia. ¹¹The Royal Women's Hospital, Parkville, Victoria 3052, Australia. ¹²Centre for Cancer Research, University of Sydney at Westmead Millennium Institute, and Department of Gynaecological Oncology, Westmead Hospital, Sydney, New South Wales 2145, Australia. ¹³Crown Princess Mary Cancer Centre and University of Sydney at Westmead Hospital, Westmead, Sydney, New South Wales 2145, Australia. ¹⁴Prince of Wales Clinical School, University of New South Wales, Sydney, New South Wales 2031, Australia. ¹⁵Victorian Institute of Forensic Medicine, Southbank, Victoria 3006, Australia. ¹⁶Victorian Life Sciences Computation Initiative, Carlton, Victoria 3053, Australia. ¹⁷La Trobe Institute for Molecular Science, Bundoora, Victoria 3083, Australia. ¹⁸Dana-Farber Cancer Institute, Boston, Massachusetts 02115-5450, USA. ¹⁹University of Chicago, Chicago, Illinois 60637, USA. ²⁰The University of Western Australia, Crawley, Western Australia 6009, Australia. ²¹Department of Biochemistry and Molecular Biology, University of Melbourne, Parkville, Victoria 3052, Australia. [†]A list of authors and affiliations appears in the Supplementary Information.

*These authors contributed equally to this work.

§These authors jointly supervised this work.

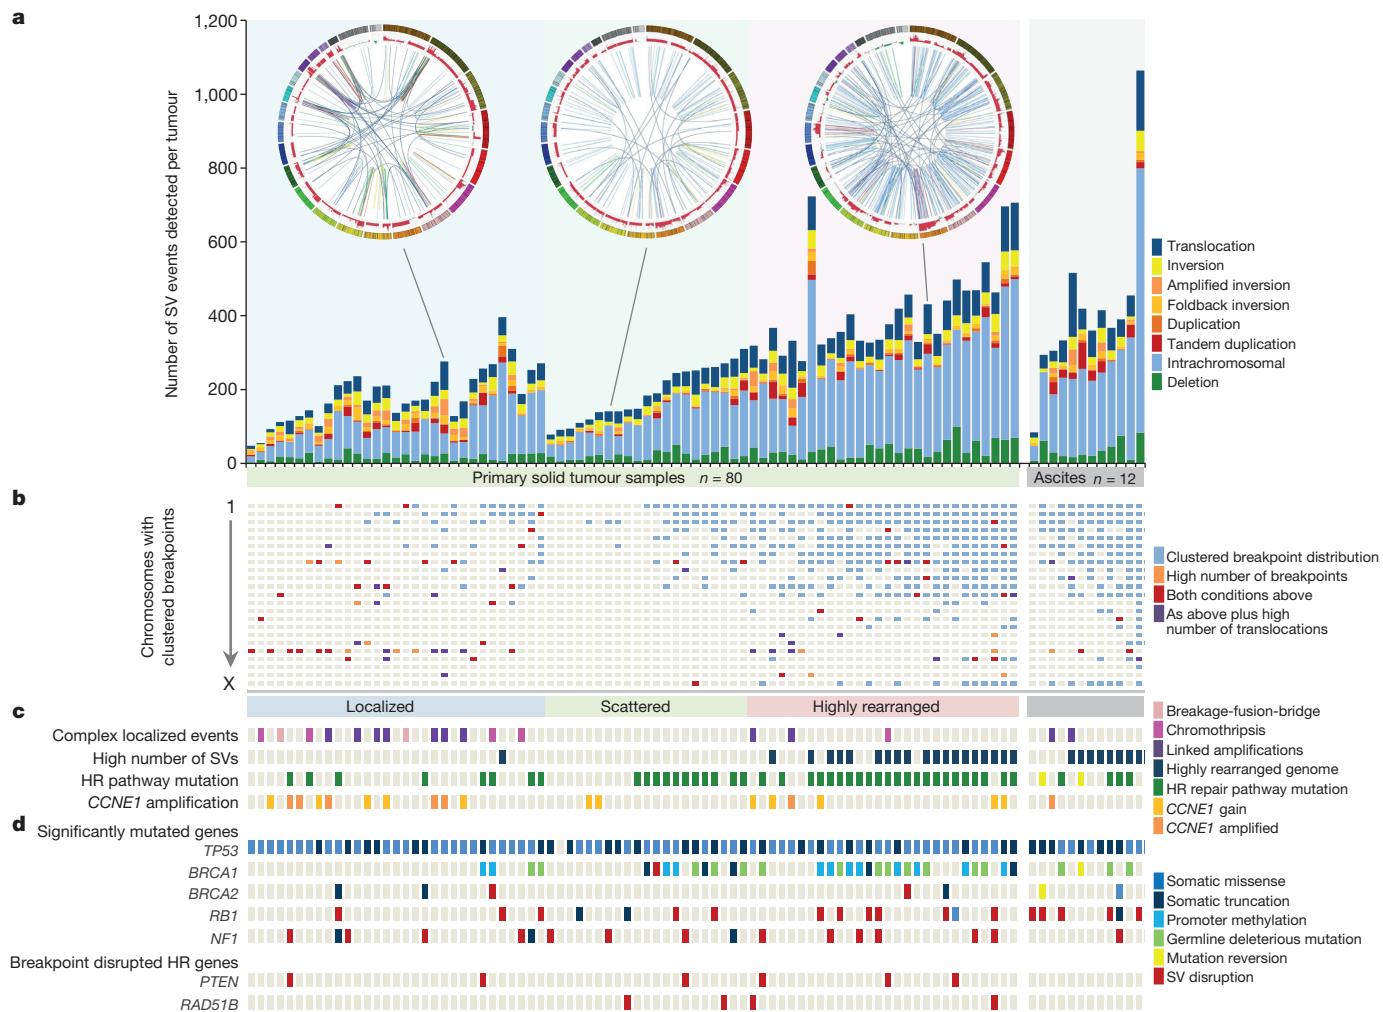


Figure 1 | Genomic features of HGSC. **a**, Structural variants (SVs) in 80 primary solid tumours and 12 ascites samples from 92 individual cases. **b**, Highly rearranged genomes have at least 7 chromosomes containing clustered breakpoint regions (blue squares), high density of breakpoints (orange), both previous conditions (red) or high density of inter-chromosomal translocations (purple). **c**, Genomes with less than 7 chromosomes with

clustered breakpoints show localized complex rearrangements (chromothripsis, pink; breakage-fusion-bridge, pale pink; and amplification linked by translocations, purple) or scattered breakpoints. HR, homologous recombination. **d**, Significantly mutated driver genes (Intogen q -value <0.01) and homologous recombination repair genes disrupted by rearrangements.

Supplementary Table 4). CCNE1 gain/amplification (19% cases) was largely exclusive of BRCA1/2 pathway disruption (Fig. 1).

A total of 36,561 somatic structural variants were detected in primary and recurrence samples, ranging from 48 to 1,064 per tumour (Fig. 1a and Supplementary Table 5). The frequency of detected structural variant events (Extended Data Fig. 2a) or single nucleotide variants (SNVs) (Extended Data Fig. 2b) was not related to tumour sample purity. We observed cases with highly rearranged genomes, a low number of scattered breakpoints or a high concentration of events involving one or few chromosomes (Fig. 1a, b and Supplementary Table 6). Similar to basal-like breast cancers¹⁰, primary solid tumour samples with BRCA1 inactivation had more structural variants ($P < 0.0001$; Extended Data Fig. 2c), involving numerous chromosomal regions (Fig. 1a–d). Tumours with BRCA1 mutations also had a higher number of genome-wide ($P < 0.001$) and coding ($P < 0.001$) small mutations (SNVs, and insertions or deletions (indels)) compared to homologous-recombination-intact samples (Extended Data Fig. 2d, e). BRCA1 mutation is frequently associated with a pronounced intratumoural immune response^{11–13}, which may relate to the greater number of mutations observed in these tumours. We identified one instance of BRCA2 disruption by chromothripsis¹⁴, but this mutational mechanism was infrequent, as was the occurrence of breakage-fusion-bridge cycles¹⁵ (Fig. 1c, Extended Data Fig. 3 and Supplementary Information).

Additional highly complex focal events were observed, including a pattern of amplified loci linked by high-density inter-chromosomal translocations similar to chromoplexy¹⁶, and chromothriptic-associated production of double minute chromosomes^{17,18} (Supplementary Information). Chromosome 19 commonly contained clusters of localized breakpoints, particularly samples with amplification involving 19q12 and the CCNE1 locus (Fig. 1b, c).

Previous exome analyses concluded that apart from TP53, somatic point mutations in driver genes are infrequent in primary HGSC³, an observation consistent with our findings (Fig. 1d and Supplementary Table 7). However, as exome studies have a limited ability to detect gene mutation by structural rearrangement, we considered inactivation of genes by disruption of transcriptional units (gene breakage). Although NF1 and RB1 were inactivated by truncating point mutations and indels in only 6% of primary samples (5 out of 80), inclusion of gene breakage raised the frequency of inactivating mutations to 20% for NF1 and 17.5% for RB1, an observation supported by expression levels and transcriptome sequence data (Fig. 1d, Extended Data Fig. 4 and Supplementary Table 8). Most involved breakage without deletion, and therefore would not have been detected by copy analysis. Gene inactivation by breakage was also seen for PTEN and RAD51B (Extended Data Fig. 4).

We searched for gene fusion events and identified 822 genes involved in intergenic fusion events potentially capable of producing a fused transcript but without evidence of recurrent, biologically plausible driver events in primary HGSC (Supplementary Information). We did not detect any evidence supporting previously reported *ESRRA* and *C11orf20* (also known as *TEX40*)¹⁹, and *CDKN2D* and *WDFY2* (ref. 20) fusions in our cohort (Supplementary Information).

Mutational signatures

Analysis of mutations across cancer genomes can reveal distinct patterns of single nucleotide substitution, many of which are associated with carcinogen exposure or mutational processes operative during tumorigenesis²¹. For most samples, either the previously described²¹ age- or *BRCA*-associated signatures were dominant (Fig. 2 and Extended Data Fig. 5), with a close relationship between the age signature, patient age at diagnosis ($P < 0.0001$; Extended Data Fig. 6a) and *CCNE1* gene amplification ($P < 0.01$; Extended Data Fig. 6b). In a small number of cases, the APOBEC or mismatch repair signatures were apparent (Extended Data Fig. 5 and Supplementary Information). The *BRCA* mutational signature was dominant in all 34 samples with germline or somatic inactivation of *BRCA1* or *BRCA2* ($P < 0.0001$), however, this association was less consistent in the small number of patients with other homologous recombination pathway mutations (Extended Data Fig. 6c). Similarly, primary sensitive cases were found to show the highest *BRCA* signature contribution (Extended Data Fig. 6d). A group of 12 patients with a dominant *BRCA* signature, but no mutation in the homologous recombination pathway, was identified (Fig. 2). Eleven of these patients were either resistant or refractory to primary treatment, consistent with intact homologous recombination function. Therefore, although the *BRCA* signature sensitively detected germline or somatic *BRCA1/2* mutations, it could also be dominant in samples in which there was no evidence of mutational or clinical characteristics suggestive of a homologous recombination pathway defect.

Clinical associations

We compared patterns of structural variation, gene expression, miRNA and methylation in primary tumour samples in the context of response to primary chemotherapy. Consistent with our previous findings²², amplification of 19q12 involving *CCNE1* was the dominant structural variant associated with primary treatment failure, while most patients with germline or somatic mutation in *BRCA1* or *BRCA2* had a favourable response to treatment (Extended Data Fig. 7a, b and Supplementary Information). Only two refractory patients (2 out of 12, 16.7%) had evidence of homologous recombination deficiency, both involving somatic methylation of *BRCA1* (Extended Data Fig. 5), consistent with the previous observation³ that methylation is not functionally equivalent to germline mutation in mediating platinum sensitivity. Although somatic mutation frequency (SNVs and indels) was significantly lower in resistant ($P < 0.05$) and refractory ($P < 0.01$) cases compared to sensitive samples (Extended Data Fig. 7c, d), other than *CCNE1* amplification and *BRCA1/2* mutation, we did not observe notable differences in the type or frequency of driver mutations, patterns of methylation or miRNA expression in the clinical subsets (Supplementary Information). We investigated the approximately one-third of tumours with neither a discernable defect in the homologous recombination pathway, nor *CCNE1* amplification or gain. Both progression-free and overall survival in these patients were poor relative to the homologous recombination defective group, similar to *CCNE1* gained/amplified patients in both our study and The Cancer Genome Atlas (TCGA) independent data set³ (Extended Data Fig. 7b). There was no evidence for high expression of *CCNE1* (Extended Data Fig. 7e) or over representation of other driver events in this group. We previously reported²³ and validated^{3,24} four molecular subtypes of HGSC defined by distinct gene expression profiles and clinical outcome. Molecular subtype analysis of the poor-prognosis, homologous-recombination-intact/*CCNE1* wild-type tumours, showed an underrepresentation of the favourable outcome C2/immune subtype in both our series and TCGA (Extended Data Fig. 7f).

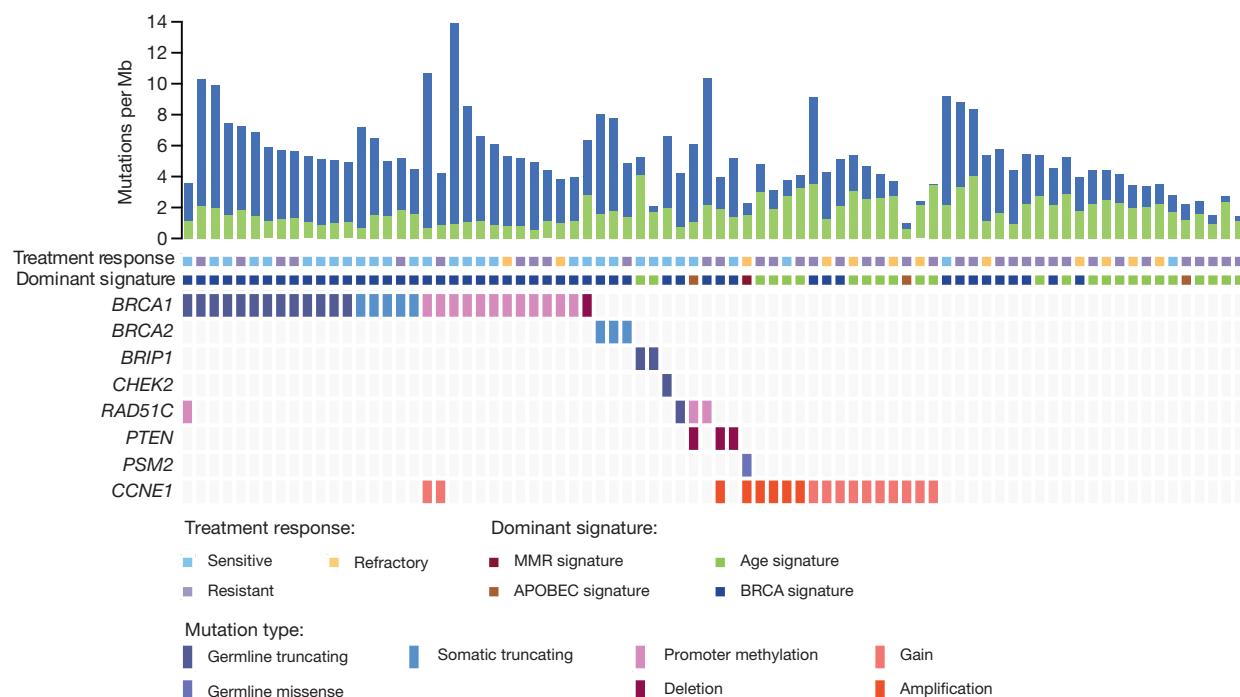


Figure 2 | Mutational signatures of primary high-grade serous ovarian cancer. Bars represent *BRCA* (blue) and age (green) signature mutation rates for 80 primary tumours. Coloured squares indicate primary response to chemotherapy and dominant mutational signature. Deleterious germline and somatic mutations in the homologous recombination pathway are shown, as

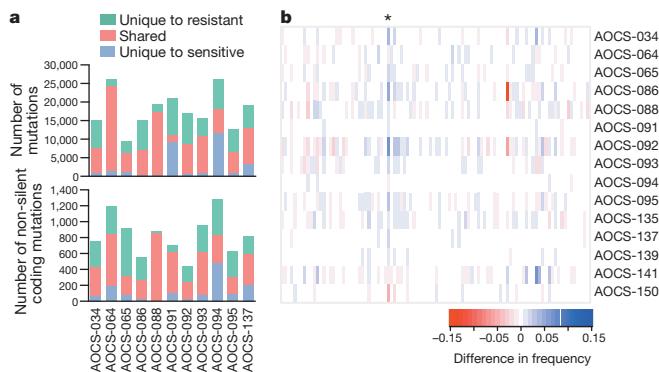
well as cases with *CCNE1* gain or amplification, and one case with a germline mismatch repair gene (*PMS2*) mutation. Samples are grouped by mutated gene and ordered left to right by descending *BRCA* signature mutation rate. Mb, megabases.

Acquired resistance

Surgery or biopsy is rarely performed in recurrent HGSC and so we relied on collection of tumour cells from ascites, drained to relieve abdominal distension, to examine relapsed disease. We focused on patients who were sensitive to initial treatment but had failed subsequent therapy (Supplementary Information). Relapse samples were found to have a higher mutational burden (SNVs and indels) at the coding and non-coding level than matched primary samples (Fig. 3a). In principle, this may reflect wider representation of tumour deposits within ascites, however, few differences were observed between primary ascites and primary tumour samples (Extended Data Fig. 8a, b).

We observed a significant relationship between the number of non-silent coding SNVs unique to the relapse samples and the number of courses of platinum-based chemotherapy the patient received ($P = 0.046$; Extended Data Fig. 8c, d), suggesting that HGSC continue to evolve during treatment. To evaluate the effect of chemotherapy on the HGSC genome, we investigated mutations that were unique to the relapse samples. We observed distinct changes in the frequency of mutations in certain trinucleotide contexts in the relapse samples, with a significant increase in C(C>T)C mutations across the 15 paired cases ($P = 0.023$) (Fig. 3b). We also observed an increase in the frequency of dinucleotide substitutions in resistant (post-treatment) samples (Supplementary Information). However, we found no evidence of platinum-induced mutations specifically occurring in driver genes. We interrogated non-silent coding mutations that were unique to the relapse samples to find those that may be involved in the development of resistance. We identified 32 genes that were mutated in more than one of the paired cases; however, none was clearly linked to chemoresistance (Extended Data Fig. 8e and Supplementary Information).

We investigated other possible mechanisms of resistance, including reversion of germline mutations in *BRCA1* or *BRCA2* (ref. 25). We identified five cases with reversion mutations (Extended Data Fig. 9a), including two patients where more than one reversion mutation was identified (Extended Data Fig. 9). A notable example involved a rapid autopsy patient with a *BRCA2* germline mutation in which five independent reversion events were detected at death (Supplementary Information). This patient was non-responsive to the PARP inhibitor olaparib and to carboplatin given after her disease recurrence (Extended Data Fig. 9c). We subsequently performed deep amplicon sequencing flanking the *BRCA2* germline mutation and reversion sites in additional metastatic deposits from this patient, and identified at least a further seven high confidence reversion events (Fig. 4a, Extended Data Fig. 9d, e and Supplementary Information). We found



no evidence of a deleter phenotype in the WGS data from this patient that might explain the propensity for reversion of the germline *BRCA2* allele (Supplementary Information). Evaluation of reversion allele frequency showed that several subclonal reversion events were detected in individual deposits, particularly in the abdomen where the presence of ascites may facilitate the spread of tumour cells²⁶ (Extended Data Fig. 9c). Reversion of mutant alleles to a wild-type sequence has been reported previously^{25,27}. Minor subclones that have reverted to a wild-type sequence would be difficult to distinguish from contaminating stromal cells and therefore additional reversion events may be present at a low level.

In the second autopsy case, there was no evidence of reversion of the *BRCA1* germline allele. The primary tumour consisted of sheets of epithelial cancer cells with limited stroma and was classified as C2/immune subtype²³ by transcriptome profiling, consistent with our previous data showing enrichment of the C2/immune subtype in *BRCA1* patients¹¹. However at autopsy, histological examination revealed an extensive desmoplastic stromal reaction (Fig. 4b) and a C1 subtype (Supplementary Information). Fibrotic stromal reactions are prominent in the poor outcome C1 molecular subtype of primary HGSC²³, and have been associated with poor drug uptake and primary chemoresistance in other solid cancers²⁸.

BRCA1 promoter methylation is an important somatic driver in ~11% of HGSCs²⁹. We identified one patient in which the primary sensitive sample showed extensive promoter methylation and low *BRCA1* expression (Fig. 4c), while the relapse sample had lost *BRCA1* methylation and the gene was expressed at comparable levels to homologous-recombination-intact tumours. Comparison of global methylation patterns of primary and recurrence samples suggested a specific rather than generalized altered methylation status at relapse in this patient (Supplementary Information).

We observed that on average there were more structural variants (~1.6 times) detected in recurrence than in primary samples (Supplementary Table 5). We therefore searched for rearrangements that were unique to the post-treatment samples and that could influence chemotherapy response. Gene breakage was identified in pro-apoptotic genes *FOXO1* and *BCL2L11*, each unique to the relapse sample of a single case (Supplementary Table 8). Breakpoints in each gene were identified in a further two primary chemoresistant samples (Supplementary Information).

We also sought instances of capture and rearrangement of promoter sequences that may drive expression of a partner pro-resistance gene. Among the 114 samples, we identified 36 events in 30 unique genes. These included two recurrent promoter fusions involving *ABCB1*, which encodes the drug efflux pump MDR1. In these tumours, an intergenic deletion of ~250 kilobases (kb) juxtaposed, head-to-tail, the promoter and non-coding exon 1 of *SLC25A40* with exon 2 of *ABCB1*, resulting in a fused transcript (Fig. 4d and Extended Data Fig. 10a). There was no evidence that these events were present in matched primary, sensitive samples. Analysis of transcriptome data showed an increase in expression of *ABCB1* in both recurrent samples with the putative fusion, and a corresponding decrease in *SLC25A40*, relative to matched sensitive samples (Extended Data Fig. 10b). Indeed, the expression of *ABCB1* in these cancers was among the highest found in the cohort (Fig. 4e). Expression of a fused transcript for one of the samples where remaining material was available was confirmed by reverse transcription PCR (RT-PCR) (Extended Data Fig. 10c). Among the transcriptome data for the relapse samples we identified five additional outliers with high *ABCB1* expression. Of these, three had translocations or insertions involving the 5' end of *ABCB1* (Extended Data Fig. 10a), which may have affected the regulation of *ABCB1*. Using primers specific for the predicted fusion, RT-PCR and Sanger sequencing, we detected a fusion product in one of these samples, suggesting the presence of an additional subclone not apparent at the coverage obtained by WGS (Supplementary Information). We evaluated an additional 51 samples collected from

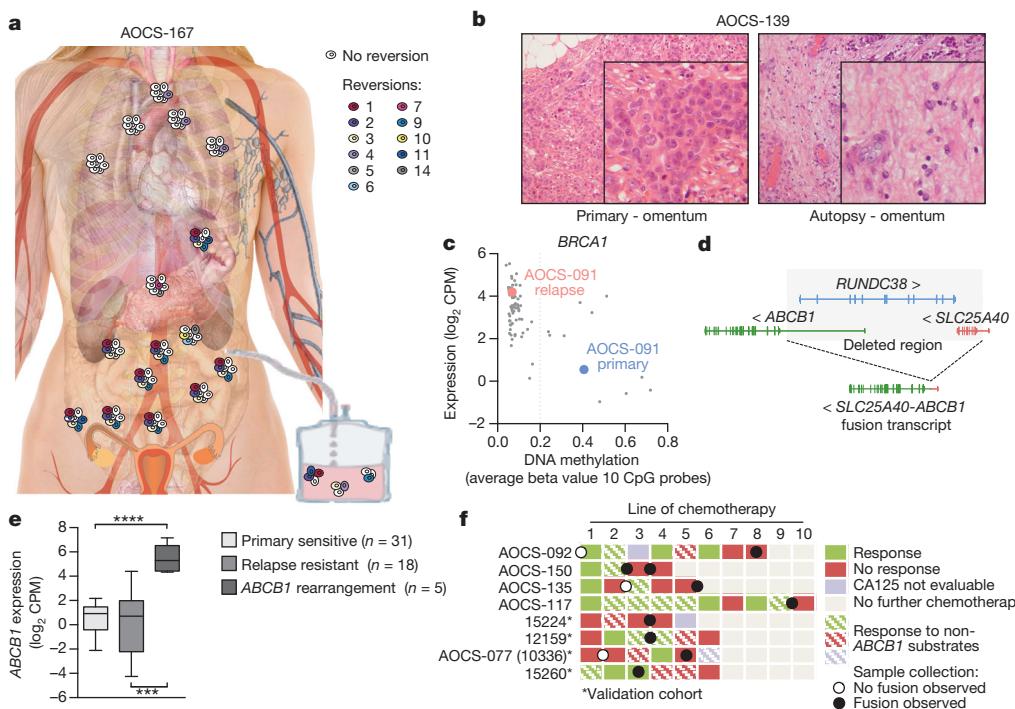


Figure 4 | Molecular changes associated with acquired resistance. **a**, In total, 11 high-confidence reversion events were identified in 16 out of 18 samples collected at autopsy in a *BRCA2* carrier. Additional large deletion event identified by WGS not shown. **b**, Haematoxylin-and-eosin-stained primary (left) and autopsy (right) tumour sections showing epithelial sheets and extensive stromal reaction, respectively. Original magnification, $\times 200$. **c**, *BRCA1* promoter methylation and low expression in AOCS-091 primary sensitive sample compared to hypomethylation and increased expression in

women with recurrent HGSC and found a further 4 independent fusion events (Extended Data Fig. 10d), all involving samples with above median *ABCB1* expression (Extended Data Fig. 10e). MDR1 (encoded by *ABCB1*) is an efflux pump for various chemotherapeutic agents used in the treatment of ovarian cancer including paclitaxel, etoposide and doxorubicin. All patients with the fusion had been exposed, and their tumours failed to respond to one or more MDR1 substrates (Fig. 4f).

Discussion

We performed whole-genome analysis of a disease that is characterized molecularly by extensive structural genomic variation and clinically by initial chemo-sensitivity, followed by the frequent emergence of chemotherapy resistance. Despite our finding that gene breakage substantially increases the frequency of recurrent mutations in *NF1* and *RB1*, few genes other than *TP53* are recurrently mutated in HGSC at a high level. The lack of driver protein-coding fusions in primary tumour samples indicates that these are less common than previously suggested^{19,20}.

While the prevailing view is that primary HGSC lacks recurrent actionable point mutations, the situation in relapse disease is unknown. We did not detect recurrent point mutations that are currently actionable in relapse samples, suggesting that, at best, only low frequency events are likely to be uncovered using personalized genomic evaluation of patients with recurrent HGSC. A recent analysis of DNA changes in *Caenorhabditis elegans* exposed to platinum recorded changes in pyrimidine-rich regions³⁰, similar to those we observed here in post-treatment samples. Although there seems to be a chemotherapy imprint on the tumour genome that includes non-silent coding changes, these did not result in mutation of known driver or resistance genes in our series.

matched resistant relapse sample. Grey dots represent all other primary tumours ($n = 80$). **d**, Schematic of deletion and *SLC25A40-ABCB1* predicted transcript. **e**, Box plots of *ABCB1* expression; lines indicate median and whiskers show range (**P < 0.001, ****P < 0.0001, two-sided t-test). **f**, Treatment response in patients with *SLC25A40-ABCB1* fusion. Primary treatment (line 1) includes both surgery and chemotherapy, and CA125 response cannot distinguish between effects of the two treatments.

We observed a range of molecular changes associated with acquired chemoresistance. Reversion of *BRCA1* and *BRCA2* germline alleles has been reported previously^{25,31,32}, and we extend these findings with a report of an apparent reversion of a methylated *BRCA1* allele in one patient. Analogous to convergent mutation of *SETD2* in renal cancer³³, we found multiple reversion events in two patients, including one at death with 12 high confidence independent *BRCA2* reversion events, underscoring a daunting potential for evolution of resistance in HGSC genomes. We note that this patient did not have debulking surgery and therefore may have had a large tumour pool from which reverted clones could arise under selection. It is of interest to determine whether the use of neoadjuvant chemotherapy³⁴ in mutation carriers increases the risk of emergence of reverted germline alleles. A surprising feature of this case was that multiple reversion events were found in individual tumour deposits, whereas distinct foci with separate, largely monoclonal events may have been expected. Furthermore, clones were shared between metastatic sites, indicating metastasis to metastasis seeding and possible cooperative interactions between clones. The multiple reversions seen in one end-stage patient contrasted with no *BRCA1* reversion in another patient, where an extensive desmoplastic stromal reaction was observed at autopsy. Tumour desmoplasia has been associated with poor drug uptake and chemoresistance in pancreatic cancer²⁸, and it is possible that the stromal reaction in this patient allowed tumour survival despite apparent retention of a homologous-recombination-defective phenotype.

The *ABCB1* gene, encoding the multidrug-resistant protein 1 (MDR1), mediates rapid efflux of many chemotherapeutic agents including paclitaxel³⁵, which forms part of standard primary chemotherapy in HGSC. We found up-regulation of this transporter apparently through promoter fusion and translocation involving the 5' region of the gene, in approximately 8% of HGSC recurrence samples.

Inhibition of MDR1 activity has not proven to be an effective strategy so far, owing to a considerable increase in paclitaxel toxicity^{36,37}. Identification of patients with this novel fusion event may allow clinicians to prioritize treatment with chemotherapy that is not a substrate of MDR1, more targeted use of MDR1 inhibitors, and use of new PARP inhibitors that are poor MDR1 substrates³⁸.

Collectively, our findings underscore the heterogeneity and apparent adaptability of the HGSC genome under the selective pressure of chemotherapy, and indicate that overcoming resistance to conventional chemotherapy will require a diversity of approaches.

Note added in proof. The C > T mutations in pyrimidine-rich regions we observed in post-treatment samples closely resembled those associated with temozolomide treatment of glioblastoma patients³⁹.

Online Content Methods, along with any additional Extended Data display items and Source Data, are available in the online version of the paper; references unique to these sections appear only in the online paper.

Received 15 December 2014; accepted 16 March 2015.

- Vaughan, S. *et al.* Rethinking ovarian cancer: recommendations for improving outcomes. *Nature Rev. Cancer* **11**, 719–725 (2011).
- Ahmed, A. A. *et al.* Driver mutations in TP53 are ubiquitous in high grade serous carcinoma of the ovary. *J. Pathol.* **221**, 49–56 (2010).
- The Cancer Genome Atlas Research Network. Integrated genomic analysis of ovarian cancer. *Nature* **474**, 609–615 (2011).
- Gorringe, K. L. *et al.* High-resolution single nucleotide polymorphism array analysis of epithelial ovarian cancer reveals numerous microdeletions and amplifications. *Clin. Cancer Res.* **13**, 4731–4739 (2007).
- Ciriello, G. *et al.* Emerging landscape of oncogenic signatures across human cancers. *Nature Genet.* **45**, 1127–1133 (2013).
- Mukhopadhyay, A. *et al.* Development of a functional assay for homologous recombination status in primary cultures of epithelial ovarian tumor and correlation with sensitivity to poly(ADP-ribose) polymerase inhibitors. *Clin. Cancer Res.* **16**, 2344–2351 (2010).
- Bashashati, A. *et al.* Distinct evolutionary trajectories of primary high-grade serous ovarian cancers revealed through spatial mutational profiling. *J. Pathol.* **231**, 21–34 (2013).
- Castellarin, M. *et al.* Clonal evolution of high-grade serous ovarian carcinoma from primary to recurrent disease. *J. Pathol.* **229**, 515–524 (2013).
- Cowin, P. A. *et al.* LRP1B deletion in high-grade serous ovarian cancers is associated with acquired chemotherapy resistance to liposomal doxorubicin. *Cancer Res.* **72**, 4060–4073 (2012).
- Popova, T. *et al.* Ploidy and large-scale genomic instability consistently identify basal-like breast carcinomas with BRCA1/2 inactivation. *Cancer Res.* **72**, 5454–5462 (2012).
- George, J. *et al.* Nonequivalent gene expression and copy number alterations in high-grade serous ovarian cancers with BRCA1 and BRCA2 mutations. *Clin. Cancer Res.* **19**, 3474–3484 (2013).
- Clarke, B. *et al.* Intraepithelial T cells and prognosis in ovarian carcinoma: novel associations with stage, tumor type, and BRCA1 loss. *Mod. Pathol.* **22**, 393–402 (2009).
- Soslow, R. A. *et al.* Morphologic patterns associated with BRCA1 and BRCA2 genotype in ovarian carcinoma. *Mod. Pathol.* **25**, 625–636 (2012).
- Stephens, P. J. *et al.* Massive genomic rearrangement acquired in a single catastrophic event during cancer development. *Cell* **144**, 27–40 (2011).
- Kinsella, M. & Bafna, V. Combinatorics of the breakage-fusion-bridge mechanism. *J. Comput. Biol.* **19**, 662–678 (2012).
- Baca, S. C. *et al.* Punctuated evolution of prostate cancer genomes. *Cell* **153**, 666–677 (2013).
- Rausch, T. *et al.* Genome sequencing of pediatric medulloblastoma links catastrophic DNA rearrangements with TP53 mutations. *Cell* **148**, 59–71 (2012).
- Nones, K. *et al.* Genome-wide DNA methylation patterns in pancreatic ductal adenocarcinoma reveal epigenetic deregulation of SLIT-ROBO, ITGA2 and MET signaling. *Int. J. Cancer* **135**, 1110–1118 (2014).
- Salzman, J. *et al.* ESRRA-C11orf20 is a recurrent gene fusion in serous ovarian carcinoma. *PLoS Biol.* **9**, e1001156 (2011).
- Kannan, K. *et al.* CDKN2D-WDFY2 is a cancer-specific fusion gene recurrent in high-grade serous ovarian carcinoma. *PLoS Genet.* **10**, e1004216 (2014).
- Alexandrov, L. B. *et al.* Signatures of mutational processes in human cancer. *Nature* **500**, 415–421 (2013).
- Etemadmoghadam, D. *et al.* Integrated genome-wide DNA copy number and expression analysis identifies distinct mechanisms of primary chemoresistance in ovarian carcinomas. *Clin. Cancer Res.* **15**, 1417–1427 (2009).
- Tothill, R. W. *et al.* Novel molecular subtypes of serous and endometrioid ovarian cancer linked to clinical outcome. *Clin. Cancer Res.* **14**, 5198–5208 (2008).
- Verhaak, R. G. *et al.* Prognostically relevant gene signatures of high-grade serous ovarian carcinoma. *J. Clin. Invest.* **123**, 517–525 (2013).
- Norquist, B. *et al.* Secondary somatic mutations restoring BRCA1/2 predict chemotherapy resistance in hereditary ovarian carcinomas. *J. Clin. Oncol.* **29**, 3008–3015 (2011).
- Lengyel, E. Ovarian cancer development and metastasis. *Am. J. Pathol.* **177**, 1053–1064 (2010).
- Swisher, E. M. *et al.* Secondary BRCA1 mutations in BRCA1-mutated ovarian carcinomas with platinum resistance. *Cancer Res.* **68**, 2581–2586 (2008).
- Olive, K. P. *et al.* Inhibition of Hedgehog signaling enhances delivery of chemotherapy in a mouse model of pancreatic cancer. *Science* **324**, 1457–1461 (2009).
- Geisler, J. P., Hatterman-Zogg, M. A., Rathe, J. A. & Buller, R. E. Frequency of BRCA1 dysfunction in ovarian cancer. *J. Natl. Cancer Inst.* **94**, 61–67 (2002).
- Meier, B. *et al.* *C. elegans* whole-genome sequencing reveals mutational signatures related to carcinogens and DNA repair deficiency. *Genome Res.* **24**, 1624–1636 (2014).
- Edwards, S. L. *et al.* Resistance to therapy caused by intragenic deletion in BRCA2. *Nature* **451**, 1111–1115 (2008).
- Sakai, W. *et al.* Secondary mutations as a mechanism of cisplatin resistance in BRCA2-mutated cancers. *Nature* **451**, 1116–1120 (2008).
- Gerlinger, M. *et al.* Intratumor heterogeneity and branched evolution revealed by multiregion sequencing. *N. Engl. J. Med.* **366**, 883–892 (2012).
- Vergote, I. *et al.* Neoadjuvant chemotherapy or primary surgery in stage IIIC or IV ovarian cancer. *N. Engl. J. Med.* **363**, 943–953 (2010).
- Murray, S., Briassoulis, E., Linardou, H., Bafaloukos, D. & Papadimitriou, C. Taxane resistance in breast cancer: mechanisms, predictive biomarkers and circumvention strategies. *Cancer Treat. Rev.* **38**, 890–903 (2012).
- Binkhathlan, Z. & Lavasanifar, A. P-glycoprotein inhibition as a therapeutic approach for overcoming multidrug resistance in cancer: current status and future perspectives. *Curr. Cancer Drug Targets* **13**, 326–346 (2013).
- Lhomé, C. *et al.* Phase III study of valsparodar (PSC 833) combined with paclitaxel and carboplatin compared with paclitaxel and carboplatin alone in patients with stage IV or suboptimally debulked stage III epithelial ovarian cancer or primary peritoneal cancer. *J. Clin. Oncol.* **26**, 2674–2682 (2008).
- Jaspers, J. E. *et al.* Loss of 53BP1 causes PARP inhibitor resistance in Brca1-mutated mouse mammary tumors. *Cancer Discov.* **3**, 68–81 (2013).
- Johnson, B. E. *et al.* Mutational analysis reveals the origin and therapy-driven evolution of recurrent glioma. *Science* **343**, 189–193 (2014).

Supplementary Information is available in the online version of the paper.

Acknowledgements The AOCS gratefully acknowledge the cooperation of the participating institutions in Australia, and also acknowledge the contribution of the study nurses, research assistants and all clinical and scientific collaborators including L. Galletta, C. Emmanuel, L. Bowes and J. Hallo. The authors acknowledge assistance from C. Anderson and D. Gwynne. The CASCADE investigators would like to thank the CASCADE Management Committee, all staff at the Victorian Institute of Forensic Medicine, D. Stevens and Tobin Brothers Funerals. The investigators would like to thank the Australia New Zealand Gynaecological Oncology Group (ANZGOG) and the women, and their families, who participated in these research programs. This work was supported by the National Health and Medical Research Council of Australia (NHMRC ID631701), Worldwide Cancer Research (09-0676) and Cancer Australia (1004673). The Australian Ovarian Cancer Study was supported by the US Army Medical Research and Materiel Command under DAMD17-01-1-0729, The Cancer Council Victoria, Queensland Cancer Fund, The Cancer Council New South Wales, The Cancer Council South Australia, The Cancer Foundation of Western Australia, The Cancer Council Tasmania and the National Health and Medical Research Council of Australia (NHMRC; ID400413, ID400281). The AOCS gratefully acknowledges additional support from S. Boldeman, the Agar family, Ovarian Cancer Australia and Ovarian Cancer Action (UK). The Gynaecological Oncology Biobank at Westmead, a member of the Australasian Biospecimen Network-Oncology group, was supported by grants from the NHMRC (ID 310670, ID 628903) and the Cancer Institute of New South Wales. The CASCADE study was supported by the Peter MacCallum Cancer Centre Foundation, and in kind by the Victorian Institute of Forensic Medicine and Tobin Brothers Funerals.

Author Contributions Project supervision: J.V.P., Nicola.W., A.D.F., S.M.G., D.D.L.B. Study design: A.M.P., E.L.C., D.E., D.W.G., S.F., P.C., Nicola.W., A.D.F., S.M.G., D.D.L.B. Sample acquisition: S.F., K.A., A.O.C.S., A.H., L.M., O.M.N., C.K., J.H., Y.E.C., P.H., M.F., M.Q., J.P., S.C., P.O.B., J.L., P.W., N.T., H.T., M.S., R.D., E.S., H.G., A.J., O.M.N., E.L., A.D.F., D.D.L.B. Sample preparation: E.L.C., D.E., D.W.G., P.C., Y.E.C., P.H., C.M., J.H. Data acquisition: E.L.C., D.E., D.W.G., S.F., D.K.M., G.M.A., T.P.H., T.S., I.H., C.N., E.N., S.M., S.I., T.J.C.B., A.N.C. Performed patient autopsy: G.Y., K.S. Sequenced data management, alignment and mutation identification: A.M.P., K.N., F.N., S.K., O.H., M.A., C.L., S.W., Q.X., J.L.F., Nick.W., S.H.N., P.J.W., J.V.P., S.M.G. Genome informatics, software tool development: A.M.P., P.J.B., K.S.K., F.N., S.K., B.P., O.H., M.A., C.L., S.W., D.F.T., Q.X., J.L.F., Nick.W., J.V.P., Nicola.W. Data analysis: A.M.P., E.L.C., D.E., D.W.G., J.G., K.N., P.J.B., K.S.K., M.C.J.Q., K.Q., C.W.B., E.C., H.S.L., G.A.Y., W.A., R.W.T., A.L., N.H., R.B., J.E., M.D., R.V., C.S., J.V.P., Nicola.W., S.M.G. Perform molecular/verification analysis: E.L.C., K.N., S.M., E.M. Wrote the manuscript: A.M.P., E.L.C., D.E., D.W.G., Nicola.W., A.D.F., S.M.G., D.D.L.B.

Author Information The whole genome and transcriptome sequencing data have been deposited in the European Genome-phenome Archive (EGA) repository under the accession code EGAD00001000877. Genotyping, methylation and miRNA data sets have been submitted into the Gene Expression Omnibus (GEO) accession GSE65821. A complete list of the AOCS Study Group can be found at <http://www.aocstudy.org>. Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to D.D.L.B. (david.bowtell@petermac.org).

METHODS

Patient cohort description. The study population consisted of women diagnosed with epithelial ovarian, primary peritoneal or fallopian tube cancer between 1992 and 2012. The women were treated at hospitals across Australia and were recruited through the Australian Ovarian Cancer Study (AOCS)²³ or through the Gynaecological Oncology Biobank at Westmead Hospital in Sydney. Four primary refractory cases were obtained from the Hammersmith Hospital Imperial College (London, UK) and the University of Chicago (Chicago, USA). Ethics board approval was obtained at all institutions for patient recruitment, sample collection and research studies. Written informed consent was obtained from all participants in this study.

All patients were diagnosed with serous carcinoma of high-grade (grade 2 or grade 3) and advanced stage (FIGO stage III or IV, International Federation of Gynaecology and Obstetrics) and received platinum-based chemotherapy as part of primary treatment. In total, 81 out of 92 (88%) patients also received a taxane at first line treatment; 86 out of 92 (93%) patients received a taxane during the course of their clinical care.

From previously published studies we estimated that with 92 unique patients, we would expect to achieve 90% power for 90% of genes to detect mutations that occur at a frequency of ~10% above the background rate for ovarian cancer (assuming a somatic mutation frequency of <2 per Mb)⁴⁰.

Additionally, as normal controls, 7 fallopian tube samples were obtained from Dana-Farber Cancer Institute (Boston, USA). Five samples were fallopian tube secretory epithelial cells isolated from dissociated fallopian tube fimbriae and expanded in short-term culture in WIT medium (Stemgent), two samples were fallopian tube epithelium collected from dissociated fallopian tube fimbriae without expansion *in vitro*.

The validation cohort of relapse ascites samples were collected from patients recruited through the AOCS with high-grade and advanced stage disease. In total, 50 out of 51 (98%) patients received platinum-based chemotherapy as part of primary treatment; 46 out of 51 (90.2%) patients also received a taxane at first line treatment; 47 out of 51 (92.2%) received a taxane during the course of treatment.

Clinical definitions. Progression-free survival was the time interval between histological or cytological diagnosis and disease progression based on GCIG (Gynaecological Cancer Inter Group) CA125 criteria⁴¹, imaging or clinical evaluation. Overall survival was defined as the interval between histological or cytological diagnosis and death from any cause. A response to second and subsequent lines of chemotherapy was defined as at least a 50% decrease in CA125 from an increased pre-treatment level, confirmed and maintained for at least 28 days⁴¹. A complete response to second and subsequent lines of chemotherapy was defined as normalization of CA125 from an increased pre-treatment level, confirmed and maintained for at least 28 days.

Inclusion criteria for treatment response groups. *Primary resistant and refractory:* This group included patients with evidence of disease progression within ~6 months from the end of primary treatment. An additional level of stringency was added by only including patients that demonstrated no response to subsequent lines of chemotherapy. A subset were designated as refractory, which included patients with disease progression while on primary treatment, or within one month of the end of primary treatment⁴², evidenced by persistently elevated CA125 or clinical progression.

Primary sensitive: Patients with no evidence of disease progression within 6 months of the end of primary treatment or patients that demonstrated responses, either normalization of CA125 or a 50% decrease in CA125 (from an increased pre-treatment level), to multiple lines of platinum-based chemotherapy.

Acquired resistance: Patients that failed to respond to chemotherapy for relapsed disease having previously demonstrated sensitivity to earlier lines of chemotherapy. Initial sensitivity was inferred if patients with macroscopic residual disease after initial surgery had no progression within 6 months from the end of primary treatment and/or had a complete or 50% response to chemotherapy for second and subsequent lines of treatment, before the development of resistance (that is, no response to later lines of chemotherapy) and collection of relapse sample.

Nucleic acid isolation. DNA was isolated from peripheral lymphocytes or lymphoblastoid cell lines using the salting out method for normal DNA. For tumour DNA and RNA isolation, frozen tissue sections were cryosectioned, with a section for haematoxylin and eosin staining being taken before and after serial sectioning for assessment of tumour content. For samples containing >70% tumour, 4 × 50 µm and 12 × 100 µm whole frozen tissue sections were used for DNA and RNA extractions, respectively. For samples containing <70% tumour, needle dissection of tumour cells was performed on up to 100 × 10 µm frozen sections for DNA and RNA extractions. Tumour cells were isolated from ascites using Dynabeads Epithelial Enrich (Invitrogen/Life Technologies), followed by DNA and RNA extractions. DNA extractions were performed using the DNeasy blood and tissue kit (QIAGEN). DNA was quantified using the Qubit dsDNA BR assay.

RNA was isolated using the mirVana miRNA Isolation kit (Ambion/Life Technologies). RNA quality was assessed using the RNA 6000 Nano kit on the BioAnalyzer (Agilent) and the NanoDrop was used to assess quality and determine RNA quantity.

SNP arrays and copy number analysis. Tumour and matched normal DNA was assayed with the Omni 2.5-8, V1.0 and V1.1 IlluminaBeadChips as per manufacturer's instructions (Illumina). Single nucleotide polymorphism (SNP) arrays were scanned on an iScan (Illumina), data was processed using the Genotyping module (v.1.9.4) in GenomeStudio v.2011.1 (Illumina) to calculate B-allele frequencies (BAF) and logR ratios. Genome alteration print⁴³ was used to call somatic regions of copy number change after low-quality probes assessed in the matched normal sample with GenCall (GC) score of <0.7 were removed. Commonly affected genes in regions with significant ($q < 0.05$) gain and loss were determined using GISTIC v2.0 (ref. 44) and genes within regions of copy number change were annotated according to the Ensembl v70 gene model.

Whole-genome library and sequence generation. Tumour cellularity was assessed using SNP array data and qPure⁴⁵ indicating a median tumour content of 85.3% (range 49.7–99.8). Matching germline DNA from blood mononuclear cells or lymphoblastoid cell lines were analysed for all patients. Sequence libraries were generated from 1 µg of genomic DNA using the standard library preparation technique as in the protocol in the TruSeq DNA PCR-Free Sample Preparation Guide for the TruSeq DNA Sample Preparation Kit (FC-121-2001, Illumina). Sequencing was carried out by the Illumina Genome Network to a minimum average of 30-fold base coverage for germline samples and 40-fold coverage for tumour samples.

WGS data processing and quality control. Each lane of sequencing data underwent alignment to the Genome Reference Consortium human genome assembly (GRCh37) using multi-threaded BWA⁴⁶ 0.6.2-nt resulting in sorted lane level files in sequence alignment/mapping (SAM) format then compressed and converted to binary file (BAM) created by Samtools⁴⁷ 0.1.19. Whole sample level merged BAMs, one each for matched germline and tumour samples were produced by in-house tools and optical duplicate reads marked using Picard MarkDuplicates 1.97 <http://picard.sourceforge.net>. Quality assessment and coverage estimation was carried out by qProfil and qCoverage (<http://sourceforge.net/projects/adamajava>). To test for the presence of sample or data swaps all sequence data were assessed for concordance at approximately 1.4 million polymorphic genomic positions including the genotyping array data by qSignature.

Transcriptome sequence data generation. Between 500 ng and 2 µg of total RNA (RNA integrity number >7, except for autopsy samples) was used for library preparation using the TruSeq RNA Sample Preparation v2 kit (Illumina) as per the manufacturer's instructions using the Low-Throughput protocol. All libraries were sequenced as paired-end 100-bp on a HiSeq2000 instrument (Illumina).

Transcriptome sequence data processing and quality control. Each barcode-separated lane of sequencing data underwent three streams of processing. To assess quality each was aligned to the Genome Reference Consortium human genome assembly (GRCh37) using multi-threaded BWA⁴⁶ 0.6.2-nt. RNA-seQC⁴⁸ was used to investigate RNA sequencing quality and concordance to all other sequencing was assessed at approximately 1.4 million polymorphic genomic positions by qSignature (<http://sourceforge.net/projects/adamajava>). Estimation of gene and transcript abundance was carried out using RSEM⁴⁹ for the Ensembl v70 gene and transcript model. The 'expected_count' values were transformed to \log_2 -counts per million ($\log_2\text{CPM}$) and the mean-variance relationship estimated to compute appropriate observational-level weights using the methods available in R-package Limma (<http://genomebiology.com/2014/15/2/R29>). A quantile normalization step was applied to normalize between samples before downstream analysis using the methods available in R-package Limma. A further gapped alignment and analysis with Mapsplice⁵⁰ was also carried out in order to identify insertions and deletions as well as differential splicing patterns and gene fusion transcripts.

MicroRNA expression analysis. MicroRNA expression levels were examined using the nCounterHuman v2.1 miRNA expression analysis kit (NanoString Technologies), with 100 ng total RNA as input for hybridization to reporter/capture probes. Purification was performed on the nCounter Prep Station and detection using the nCounter Digital Analyzer (NanoString Technologies), as per the manufacturer's instructions.

The nSolver Analysis Software (NanoString Technologies) was used to extract the raw counts for all samples that passed internal QC metrics, before normalization using the housekeeping genes to control for sample input.

Methylation. Methylation levels across the genome were examined on the Infinium Human Methylation 450 Bead Chip (Illumina), the samples were processed at the Australian Genome Research Facility. The data were background corrected, normalized to internal controls and quality control was performed at

the probe and sample level. A total of 572 probes failed in greater than 25% of the samples owing to their detection *P* values being greater than 0.05, and these were removed from the data⁵¹. COMBAT was used to remove batch effects⁵².

Somatic SNV variant detection. High confidence somatic substitutions were identified as the intersect of the post filtered output of the mutation detection tools qSNP⁵³ and the Unified Genotyper⁵⁴ (GATK). In addition, supporting evidence for variants was sought in matched independent and/or orthogonal data sets such as long mate pair sequencing data produced on Life Technologies SOLiD platform and HiSeq RNA transcriptome sequencing. Annotation of variants was carried out against Ensembl v.70 and consequence prediction using Ensembl perl API modules. Two methods, Mutsig⁵⁵ and IntOGen⁵⁶, were used to identify significant, recurrently mutated genes in the primary tumour data set ($n = 80$).

Somatic short insertion and deletion variant detection. Pindel⁵⁷ was used to identify indels of 1–50 bp in length and variant filtering, annotation and prioritization was carried out as previously described¹⁸. Manual review of both germline and somatic variants in genes highlighted in this study was carried out using Integrative Genomics Viewer (IGV)⁵⁸. Gene model annotation of variants was carried out against Ensembl v.70 and consequence prediction using Ensembl perl API modules for variant effect prediction across all transcripts.

Structural variant detection. Structural variants of length 50 bp or greater were identified using qSV¹⁸ (<http://sourceforge.net/projects/adamjava/>), which discovers somatic structural rearrangements from tumour and matched normal sequencing using three methodologies: discordantly mapped read pairs; clusters of soft clipping; and a split contig alignment generated from localized *de novo* assembly of unmapped plus aberrantly mapped reads. Only high confidence somatic structural variants were used in the analysis defined as events that were identified by more than one of the detection methodologies or that had at least 10 supporting reads, and where no evidence of the event was identified in the matched germline sample including low quality data.

Breakpoint annotation and consequence predictions were performed against the Ensembl v70 gene model to calculate the consequences for flanking genes or fragments of genes. Breakpoints were also integrated with the segmented copy number data to inform characterization: intrachromosomal breakpoints flanking a copy number segment of loss were annotated as deletions; those linking segments of similar DNA copy number were annotated as intrachromosomal translocations; duplications and inversions associated with increases in copy number enabled characterization of tandem duplications and amplified or foldback inversions.

Somatic variant verification. In-house developed software examined genomic positions of variants in the matched data sets scoring supporting read evidence and assessing coverage in a similar manner to methods used in a recent genomic landscape paper for cervical carcinoma⁵⁹. In addition for structural variants, SNP array data provided orthogonal evidence in support for structural variant events that cause a copy number shift when boundaries of copy number change were identified in close proximity to breakpoints.

Mutational signatures and detection of kataegis. Mutational signatures were deciphered using the framework developed by previously²¹ (<http://www.mathworks.com/matlabcentral/fileexchange/38724>). Loci of localized hyper-mutation, termed kataegis, were identified by calculating the inter-mutation distances between the genomic position of sorted, high-confidence somatic substitution variant calls for each sample. Inter-mutation distances were then segmented using piecewise constant fitting to find regions of constant inter-mutation distance⁶⁰. Parameters used for piecewise constant fitting were c525 and kmin52 (ref. 60). Putative regions of kataegis were identified as those segments containing six or more consecutive mutations with an average inter-mutation distance of less than or equal to 1,000 bp.

Detection of breakpoint clustering. Chromosomes containing highly significant non-random distribution of breakpoints were identified as recently described⁶¹ with a stringent threshold of $P < 0.00001$. Chromosomes with remarkably high numbers of rearrangement events were identified as outliers defined as a breakpoint per megabase rate exceeding 1.5 times the length of the inter-quartile range from the seventy-fifth percentile for each sample with a minimum threshold of 35 breakpoints per chromosome. Chromosomes identified to have both characteristics were inspected against criteria for chromothripsis⁶¹ or the breakage-fusion-bridge cumulative rearrangement model¹⁵. Similarly, chromosomes with high numbers of translocations were identified with a minimum threshold of 10 translocation breakpoints per chromosome following observations of localized events.

Molecular subtyping. Tumour molecular subtypes were classified using previously described expression profiles²³ as the learning set. Expression of genes between each high-grade serous molecular subtype was compared to identify subtype specific gene signatures. All genes differentially expressed at a FDR < 5% were selected as subtype specific genes. The union of all subtype specific gene sets were used as the features of the classifier. The normalized data from the learning set and transformed data from the RNA-seq data (described in detail in the previous section on transcriptome data processing) were further scaled

independently so that every gene in the two sets had zero mean and unit variance. Finally, the transformed gene expression profiles were used to identify the molecular subtypes of samples in the cohort using *k*-nearest neighbour classification algorithm. The accuracy of the classification procedure was estimated based on the samples of known molecular subtypes, as defined by previous microarray experiments²³. These samples were removed from the training set and the classification procedure repeated. Of the 19 samples with known molecular subtypes in the cohort, 17 were classified correctly (89.47% accuracy).

Verification and validation of SLC25A40-ABCB1 fusion transcript. A cohort of 51 relapse ascites samples was used to examine frequency of SLC25A40-ABCB1 fusion. RNA from ascites samples was reverse transcribed into cDNA using random primers (Promega) and M-MLV reverse transcriptase (Promega). Quantitative reverse transcription PCR was performed to examine SLC25A40 and ABCB1 expression, with GAPDH and ACTB for normalization, using the following primers: ABCB1_forward 5'-GAGAGATCCTACCAAGCGG-3', ABCB1_reverse 5'-CG AGCCTGGTAGTCAATGCT-3', SLC25A40_forward 5'-AGGAAACATTGGAT GCATTTT-3', SLC25A40_reverse 5'-AGCACTTAATTGATCATAGCAGGT-3'.

The $\Delta\Delta C_t$ method was used to calculate expression levels in ascites compared to expression in the SKOV3 cell line⁹.

Testing and verification of the presence of the SLC25A40-ABCB1 fusion transcript was performed using nested PCR with primers to exon 1 of SLC25A40 and exon 3 of ABCB1: outer_forward 5'-CGGCTCTGTGACCAAAC-3', outer_reverse 5'-TCTTGCTCCTCCATTGCGG-3'; internal_forward 5'-CCGTCACCAGGGTT ATTCC-3', internal_reverse 5'-CCCTTCAGATCATTCCGA-3'.

PCR conditions were as follows: 98°C for 30 s, 30 cycles of 98°C for 10 s, 60°C for 30 s, 72°C for 10 s, and 72°C for 10 min. PCR product from first PCR was purified using QIAquick PCR purification kit (QIAGEN), PCR product from second PCR was run on a 2% agarose gel, bands were excised and purified using Wizard SV Gel and PCR Clean-up System (Promega) before Sanger sequencing using internal primers.

Examination of BRCA2 reversion events in AOCS-167. Primers to flank the germline and reversion events identified by WGS were used to amplify the regions of interest for next-generation sequencing of additional tumour samples from AOCS-167. Each of the gene-specific primers were generated with a common sequence (CS) on the 5' end. The paired end sequence (PE-CS) primers add the Illumina adaptor sequences and a sample specific 10-bp barcode (BC) to the libraries.

BRCA2_1F 5'-TCATGATGAAAGTCTGAAGAAAATGA-3'; BRCA2_1R 5'-GAGCAAGACTCCACCTCAA-3'; BRCA2_2F 5'-AATTGTTGCAGAACATGT GAAAAGCT-3'; BRCA2_2R 5'-ACTAAACAGAGGACTTACCATGACT-3'; BRCA2_3F 5'-ACTGTTCTATGAGAAAGGTTGTGA-3'; BRCA2_3R 5'-TCCAATGTGGTCTTGCAGC-3'; BRCA2_4F 5'-AGCACATTCTACATAA ACTGTTCT-3'; BRCA2_4R 5'-TGAACCTTCAATGTGGTCTTGC-3'; CS_F 5'-ACACTGACGACATGGTCTACA-3'; CS_R 5'-TACGGTAGCAGAGACT TGGTCT-3'; PE-CS_F 5'-AATGATACGGCGACCACCGAGATCT-3'; PE-CS_R 5'-CAAGCAGAAGACGGCATACGAGAT-BC-3'.

The first PCR used the CS-gene-specific primers and Phusion High-Fidelity DNA Polymerase (Thermo Scientific) with PCR conditions as follows: 98°C for 30 s, 30 cycles of 98°C for 10 s, 60°C for 30 s, 72°C for 10 s, and 72°C for 10 min. QIAquick PCR purification (QIAGEN) was used and the product was used as a template for a second PCR using PE-CS primers with a unique barcode for each sample and PCR conditions as follows: 98°C for 10 min, 15 cycles of 98°C for 10 s, 60°C for 30 s, 72°C for 30 s, and 72°C for 10 min. Libraries were purified using AMPure XP Beads and analysed on Agilent BioAnalyser High Sensitivity DNA chip. All 18 libraries were run on a single Illumina MiSeq flowcell.

Sequencing data was adaptively trimmed (equivalent to the trimming performed by BWA with $q = 30$) then aligned to the amplicon sequences using an in-house bioinformatics tool (Peter MacCallum Cancer Centre Bioinformatics Core, unpublished). This tool utilizes the known primer sequences to align each read to its amplicon of origin, and where there is overlap in the forward and reverse a consensus sequence is generated using the highest quality base. Fastqc analysis was performed on the fastq files (<http://www.bioinformatics.babraham.ac.uk/projects/fastqc/>), the average base sequence quality was above 35 for all samples. pileup files were generated from the aligned sequence using SAMtools 0.1.18 (ref. 47) and VarScan v2.3 (ref. 62) was subsequently used to calculate the read depth at each position. The mean read depth ranged from 8,101–21,019 per sample. Indel variant calling was performed using IndelGenotyperV2 (GATK v1.0.4905)⁵⁴, with the minimum fraction of reads with the indel set to 0.001 and a minimum count of reads supporting the indel set to 1. All variants were manually reviewed in IGV.

For verification mutation specific PCR primers were designed with 1–5 bp spanning the deletion, and non-mutation specific primers used were BRCA2_2F and BRCA2_4R.

BRCA2_rev1_1R 5'-TGTGTTTCAGTGTCA-3'; *BRCA2_rev1_2R* 5'-TGTGTTTCAGTGTCA-3'; *BRCA2_rev2_1R* 5'-GGACTTACCATGACTTC-3'; *BRCA2_rev2_2R* 5'-GGACTTACCATGACTTC-3'; *BRCA2_REV3_1F* 5'-TTATACTTTAACAGGAA-3'; *BRCA2_rev4_1F* 5'-AAACATCAGGGAATGT-3'; *BRCA2_rev5_1R* 5'-GTGTTTCAGTGTCTT-3'; *BRCA2_rev5_2R* 5'-GTGTTTCAGTGTCTTCA-3'; *BRCA2_rev6_1R* 5'-ACCAGTACTTGCGCG-3'; *BRCA2_rev6_2R* 5'-ACCAGTACTTGCGAGCGAT-3'; *BRCA2_rev7_1R* 5'-TACCATGACTTGCACT-3'; *BRCA2_rev7_2R* 5'-TACCATGACTTGCACT-3'; *BRCA2_rev8_1R* 5'-GACTTGCACTCTCTTG-3'; *BRCA2_rev9_1R* 5'-CCATGACTTGCACTCT-3'; *BRCA2_rev9_2R* 5'-CCATGACTTGCACTCT-3'; *BRCA2_rev10_1F* 5'-GTTTTATACTTTAACAGCTG-3'; *BRCA2_rev11_1F* 5'-CTTTAACAGGATTGGAAT-3'; *BRCA2_rev12_1F* 5'-ACAGGATTGAAAACAT-3'; *BRCA2_rev12_2F* 5'-ACAGGATTGAAAACAT-3'; *BRCA2_rev13_1F* 5'-CAGGATTGAAAAT-3'; *BRCA2_rev14_1F* 5'-TTGGAAAAACATCAGCTG-3'; *BRCA2_rev14_2F* 5'-TTGGAAAAACATCAGCTG-3'; *BRCA2_rev15_1F* 5'-GAA AACATCAGGGAAG-3'; *BRCA2_rev15_2F* 5'-GAAAAACATCAGGGAATA-3'.

Touchdown PCR was used for mutation-specific PCR to improve sensitivity and specificity. PCR conditions were as follows: 98°C for 5 min, 10 cycles of 98°C for 10 s, 65°C for 30 s decreasing 1°C per cycle, 72°C for 10 s, 15 cycles of 98°C for 10 s, 55°C for 30 s, 72°C for 10 s and 72°C for 10 min. PCR products were run on 2% agarose gel for visual verification of presence of PCR product of correct size.

Code availability. All in house developed tools are available for download from <http://sourceforge.net/projects/adamjava>.

Statistics. No statistical methods were used to predetermine sample size, and the experiments were not randomized. Further details of statistics are provided in the Supplementary Information.

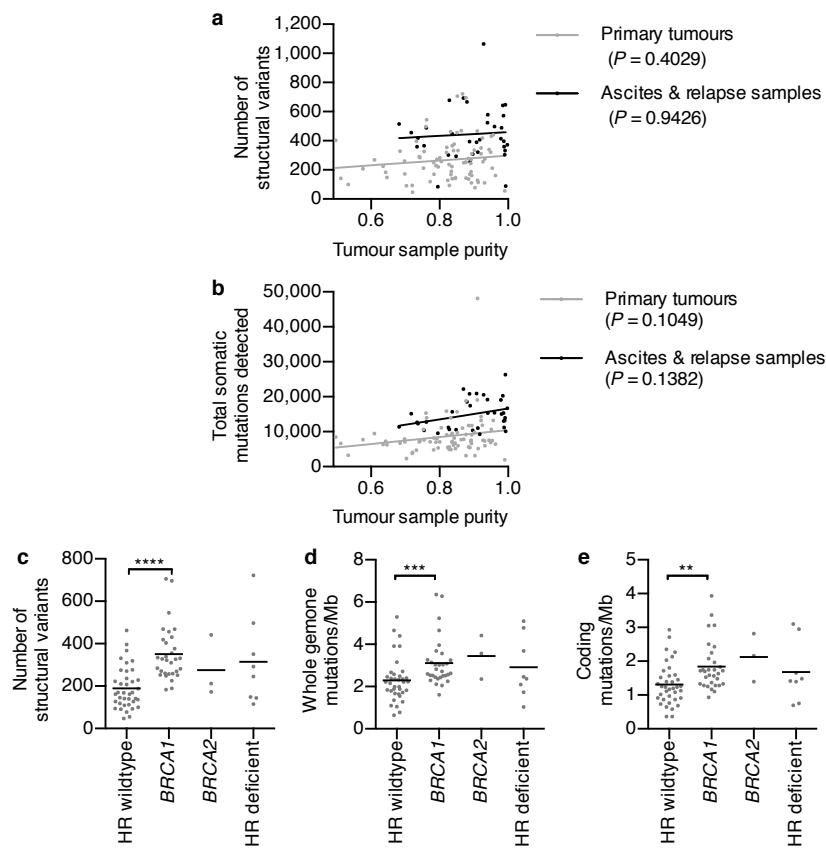
40. Lawrence, M. S. *et al.* Discovery and saturation analysis of cancer genes across 21 tumour types. *Nature* **505**, 495–501 (2014).
41. Rustin, G. J. *et al.* Definitions for response and progression in ovarian cancer clinical trials incorporating RECIST 1.1 and CA 125 agreed by the Gynecological Cancer Intergroup (GCIG). *Int. J. Gynecol. Cancer* **21**, 419–423 (2011).
42. Friedlander, M. L. *et al.* Clinical trials of palliative chemotherapy in platinum-resistant or -refractory ovarian cancer: time to think differently? *J. Clin. Oncol.* **31**, 2362 (2013).
43. Popova, T. *et al.* Genome Alteration Print (GAP): a tool to visualize and mine complex cancer genomic profiles obtained by SNP arrays. *Genome Biol.* **10**, R128 (2009).
44. Mermel, C. H. *et al.* GISTIC2.0 facilitates sensitive and confident localization of the targets of focal somatic copy-number alteration in human cancers. *Genome Biol.* **12**, R41 (2011).
45. Song, S. *et al.* qpure: A tool to estimate tumor cellularity from genome-wide single-nucleotide polymorphism profiles. *PLoS ONE* **7**, e45835 (2012).
46. Li, H. & Durbin, R. Fast and accurate short read alignment with Burrows-Wheeler transform. *Bioinformatics* **25**, 1754–1760 (2009).
47. Li, H. *et al.* The Sequence Alignment/Map format and SAMtools. *Bioinformatics* **25**, 2078–2079 (2009).
48. DeLuca, D. S. *et al.* RNA-SeQC: RNA-seq metrics for quality control and process optimization. *Bioinformatics* **28**, 1530–1532 (2012).
49. Li, B. & Dewey, C. N. RSEM: accurate transcript quantification from RNA-Seq data with or without a reference genome. *BMC Bioinformatics* **12**, 323 (2011).
50. Wang, K. *et al.* MapSplice: accurate mapping of RNA-seq reads for splice junction discovery. *Nucleic Acids Res.* **38**, e178 (2010).
51. Wilhelm-Benartzi, C. S. *et al.* Review of processing and analysis methods for DNA methylation array data. *Br. J. Cancer* **109**, 1394–1402 (2013).
52. Johnson, W. E., Li, C. & Rabinovic, A. Adjusting batch effects in microarray expression data using empirical Bayes methods. *Biostatistics* **8**, 118–127 (2007).
53. Kassahn, K. S. *et al.* Somatic point mutation calling in low cellularity tumors. *PLoS ONE* **8**, e74380 (2013).
54. McKenna, A. *et al.* The Genome Analysis Toolkit: a MapReduce framework for analyzing next-generation DNA sequencing data. *Genome Res.* **20**, 1297–1303 (2010).
55. Lawrence, M. S. *et al.* Mutational heterogeneity in cancer and the search for new cancer-associated genes. *Nature* **499**, 214–218 (2013).
56. Gundem, G. *et al.* IntOGen: integration and data mining of multidimensional oncogenomic data. *Nature Methods* **7**, 92–93 (2010).
57. Ye, K., Schulz, M. H., Long, Q., Apweiler, R. & Ning, Z. Pindel: a pattern growth approach to detect break points of large deletions and medium sized insertions from paired-end short reads. *Bioinformatics* **25**, 2865–2871 (2009).
58. Robinson, J. T. *et al.* Integrative genomics viewer. *Nature Biotechnol.* **29**, 24–26 (2011).
59. Ojesina, A. I. *et al.* Landscape of genomic alterations in cervical carcinomas. *Nature* **506**, 371–375 (2014).
60. Nik-Zainal, S. *et al.* Mutational processes molding the genomes of 21 breast cancers. *Cell* **149**, 979–993 (2012).
61. Korbel, J. O. & Campbell, P. J. Criteria for inference of chromothripsis in cancer genomes. *Cell* **152**, 1226–1236 (2013).
62. Koboldt, D. C. *et al.* VarScan 2: somatic mutation and copy number alteration discovery in cancer by exome sequencing. *Genome Res.* **22**, 568–576 (2012).

a**b**

	Primary refractory	Primary resistant	Primary chemosensitive	P value	Acquired resistant (no primary tumour)	All cases
Age						
Mean	61.7	61.3	57.6	0.188 ^a	57.7	59.7
Range	(52.5 - 74.7)	(44.5 - 78.4)	(39.2 - 73)		(45.5 - 70.8)	(39.2 - 78.4)
Primary site						
Ovary	8 (67%)	27 (73%)	28 (90%)	0.115 ^b	10 (83%)	73 (79%)
Peritoneum	4 (33%)	9 (24%)	2 (6%)		2 (17%)	17 (18%)
Fallopian tube	0 (0%)	1 (3%)	1 (3%)		0 (0%)	2 (2%)
Grade						
2	2 (17%)	4 (11%)	5 (16%)	0.745 ^b	1 (8%)	12 (13%)
3	10 (83%)	33 (89%)	26 (84%)		9 (75%)	78 (85%)
Ungraded*	0 (0%)	0 (0%)	0 (0%)		2 (17%)	2 (2%)
Stage						
III	12 (100%)	31 (84%)	25 (81%)	0.300 ^b	10 (83%)	78 (85%)
IV	0 (0%)	6 (16%)	6 (19%)		2 (17%)	14 (15%)
Residual disease						
No macroscopic	1 (8%)	2 (5%)	1 (3%)	0.876 ^b	1 (8%)	5 (5%)
Macroscopic ≤ 1 cm	6 (50%)	23 (62%)	18 (58%)		2 (17%)	49 (53%)
Macroscopic >1 cm	5 (42%)	12 (32%)	12 (39%)		4 (33%)	33 (36%)
Macroscopic (size unknown)	0 (0%)	0 (0%)	0 (0%)		5 (42%)	5 (5%)
Progression free survival						
Number of events	12 (100%)	37 (100%)	24 (77%)	<0.0001 ^c	12 (100%)	85 (92%)
Median months [^]	4.98	9.24	17.72		21.21	10.39
95% CI of median	(4.14 - 7.04)	(7.56 - 10.13)	(13.28 - 21.96)		(13.35 - 26.01)	(9.67 - 11.47)
Overall survival						
Number of events	11 (92%)	36 (97%)	19 (61%)	<0.0001 ^c	10 (83%)	76 (83%)
Median months [^]	12.92	20.22	55.86		46.04	27.11
95% CI of median	(6.41 - 34.68)	(16.31 - 22.49)	(45.96 - 69.07)		(31.76 - 64.83)	(21.50 - 34.49)
Total cases	12	37	31		12	92

Extended Data Figure 1 | Patient cohort. **a**, Summary of whole-genome sequenced patients ($n = 92$) and samples ($n = 114$). **b**, Clinical characteristics of patients by clinical response group and acquired resistant cases without matching primary tumour material. *Primary tumour ungraded; diagnosis

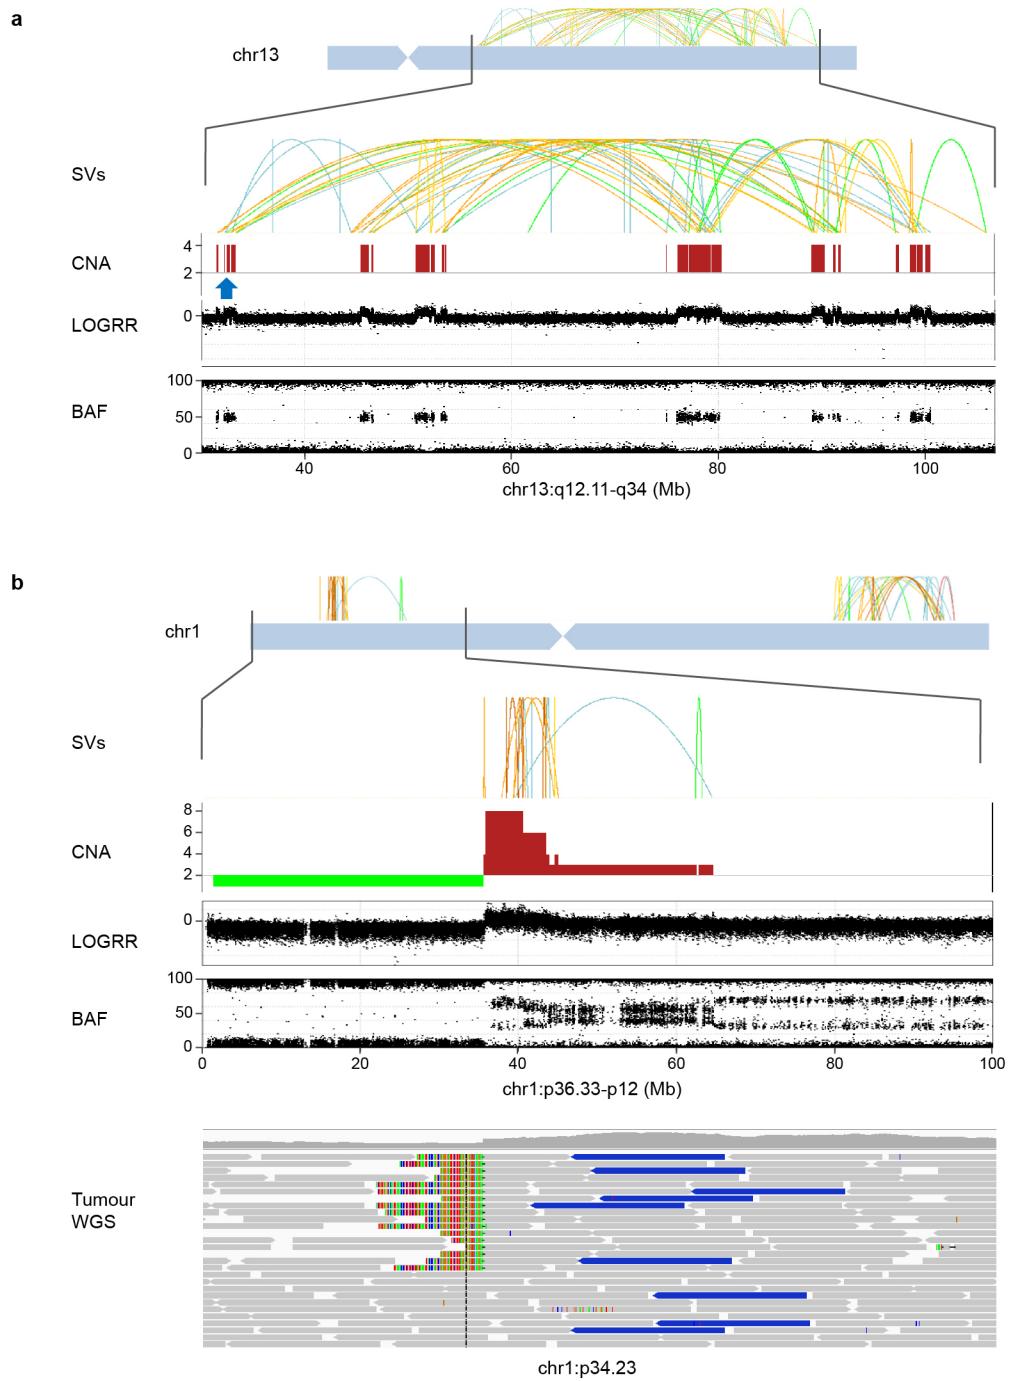
from ascites or pleural fluid. [^]Time to progression or death measured from diagnosis. ^aKruskal-Wallis, ^bFisher or ^clog-rank test P values comparing primary tumour clinical groups reported. Median follow up time of cohort was 97.3 months.



Extended Data Figure 2 | Structural variants and somatic mutations.

a, b, Scatter plot shows that the number of structural variants (a) and SNVs (b) detected per sample was not dependent on the purity of primary tumours ($n = 80$, grey) and ascites/relapse samples ($n = 34$, black). Spearman correlation P values indicated. **c**, Number of structural variants detected in primary tumours ($n = 80$) grouped by homologous recombination (HR) mutation status: HR wild-type, *BRCA1/2* altered, and HR deficient (lines

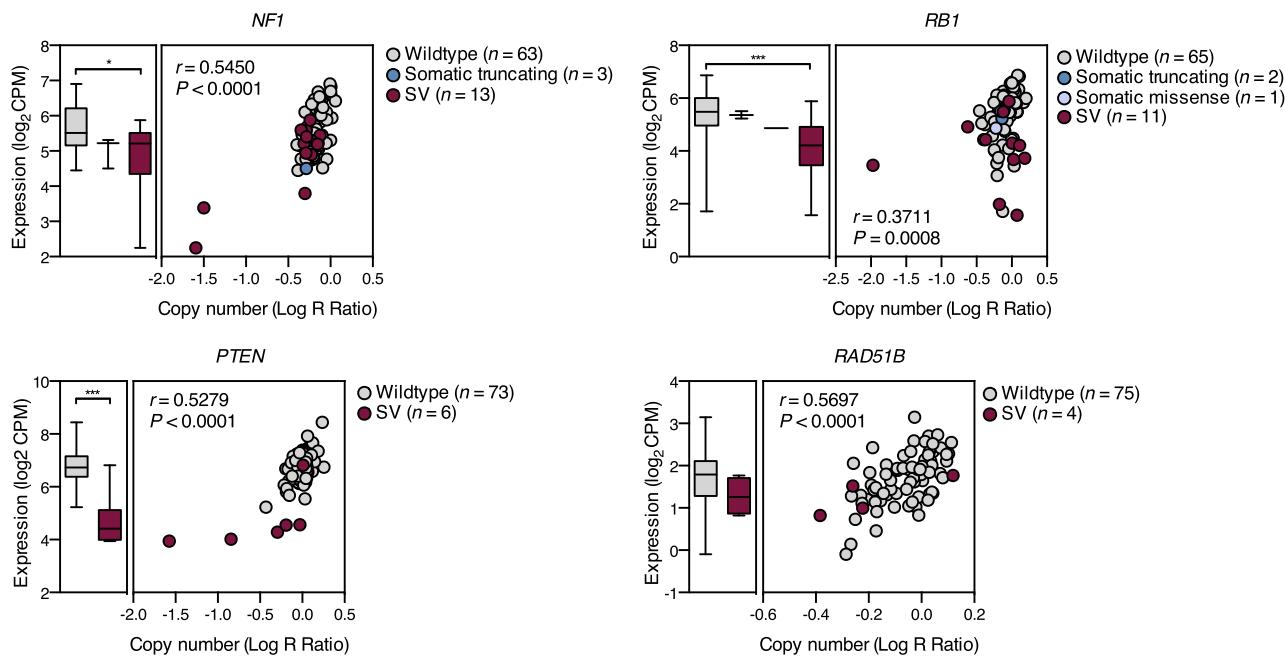
indicate mean; **** $P < 0.0001$, Kolmogorov-Smirnov test). **d, e**, Whole-genome (d) and coding mutations (e) per megabase for samples ($n = 79$) stratified by *BRCA1*, *BRCA2* or homologous recombination pathway mutation status (lines indicate mean; ** $P < 0.01$, *** $P < 0.001$, Kolmogorov-Smirnov test). Sample AOCS-166 with germline mismatch repair mutation excluded from analysis (see Supplementary Information section 4.3).



Extended Data Figure 3 | Chromothripsis and breakage-fusion-bridge.

a. Chromothripsis affecting chromosome 13, including the *BRCA2* locus (arrow), for a primary chemotherapy-resistant tumour. A high number of breakpoints (structural variants) are observed with oscillations of copy number (CNA and LOGRR) indicating regions of retained heterozygosity (BAF) for a

single haplotype. **b.** Breakage-fusion-bridge amplification is observed on chromosome 1 for a primary chemotherapy-resistant tumour. A cluster of breakpoints (structural variants), mostly inversions, on the distal p arm are associated with blocks of amplification (CNA). IGV review of the tumour WGS confirms that the telomere region has been lost.

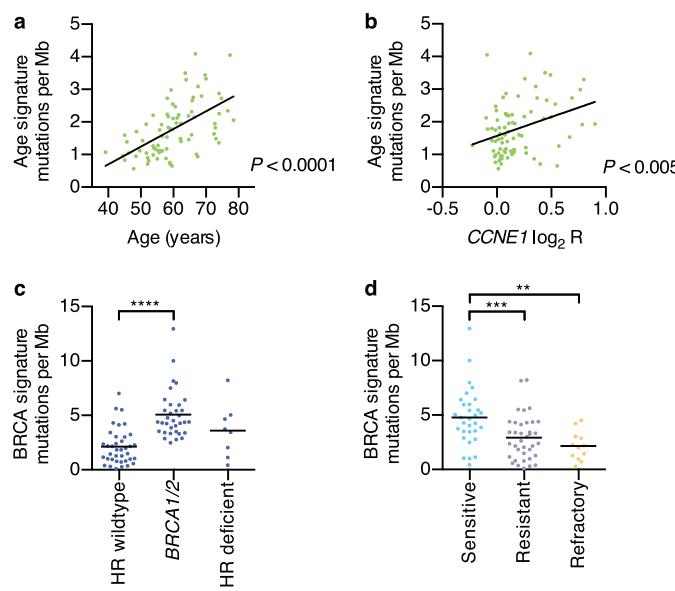


Extended Data Figure 4 | Expression of genes altered by structural variation. Scatter graphs show expression of *RB1*, *NF1*, *PTEN* and *RAD51B* plotted against copy number in primary tumours ($n = 79$; Spearman correlation analysis). Boxplots summarize expression by mutation type; lines indicate median and whiskers show range (two-tailed Mann-Whitney test; $*P < 0.05$, $***P < 0.001$). Samples with somatic interrupting structural

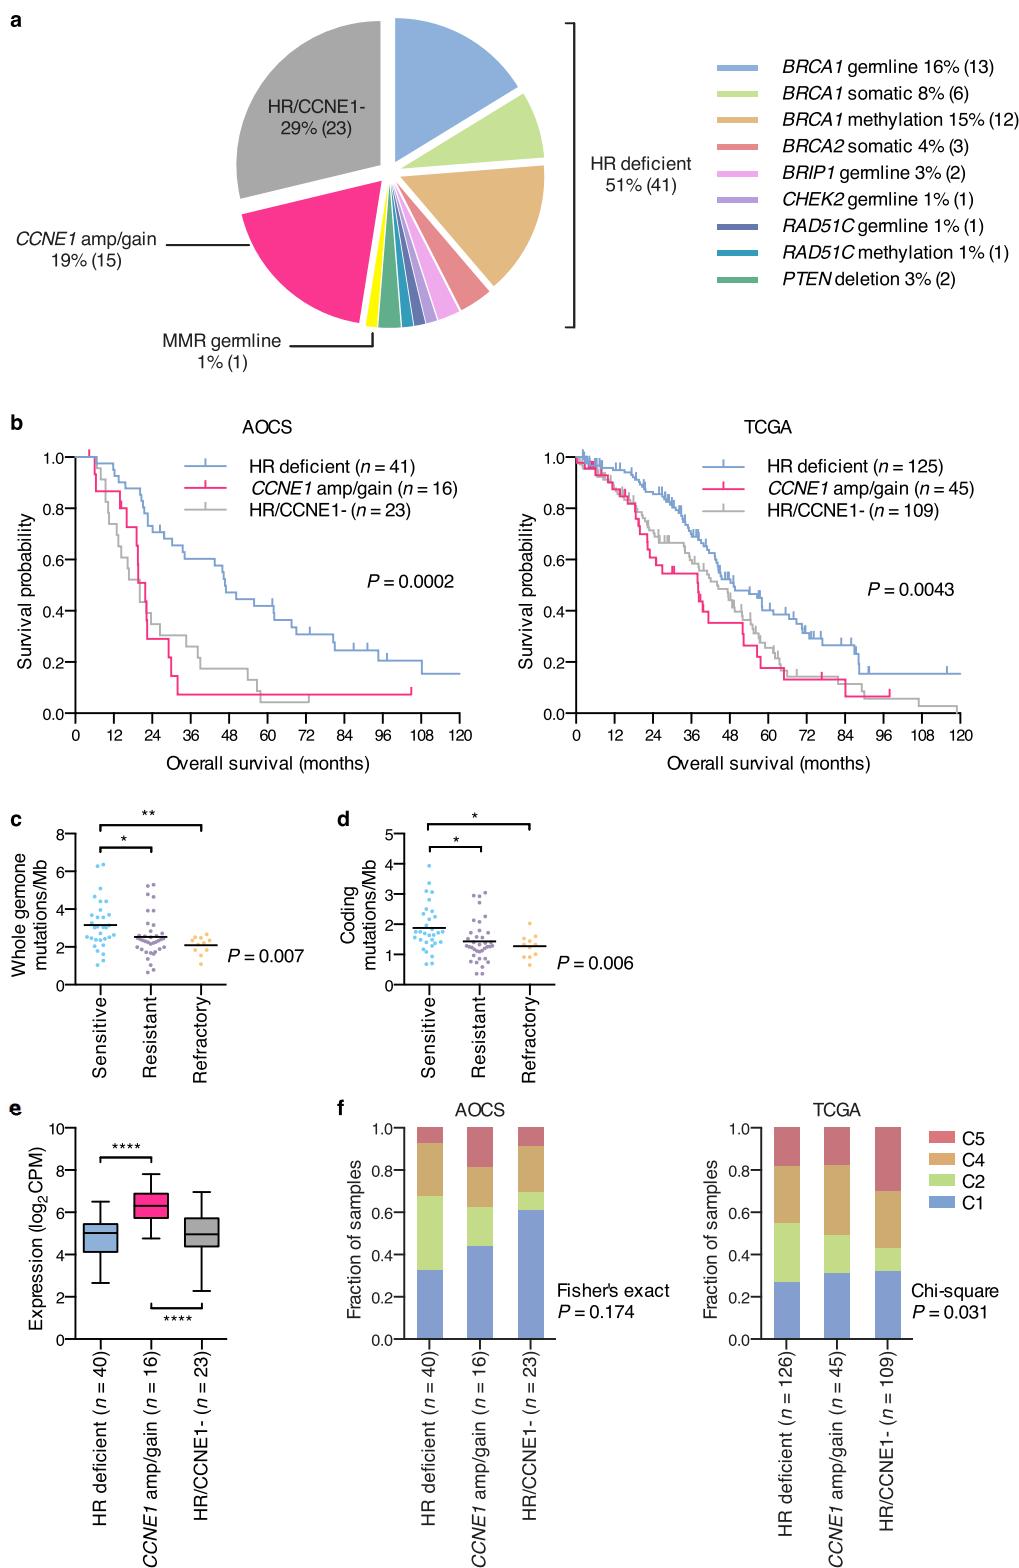
variants and truncating mutations are indicated. Not all structural variants are associated with DNA loss and translocations can be observed to produce high expression counts for non-functional transcripts. Expression values lower than the median value for tumours and controls was observed for the majority of samples with structural variant events in *RB1*, *NF1*, *PTEN* and *RAD51B*.

Study ID	Age	Response group	HR pathway status	HR gene mutation	CCNE1 status	CCNE1 log ₂ R (CBS segment)	MMR signature mut/Mb	APOBEC signature mut/Mb	Age signature mut/Mb	BRCA signature mut/Mb	Dominant signature	
AOCS-064	67.5	Sensitive	HRD	<i>BRCA1 meth</i>	-0.0	-	2.1	1.7	0.9	13.0	4	
AOCS-094	59.0	Sensitive	HRD	<i>BRCA1 meth</i>	Gain	0.40	3.8	3.0	0.7	10.0	4	
AOCS-106	64.6	Resistant	HRD	<i>RAD51C meth</i>	0.11	-	1.7	1.1	2.2	8.2	4	
AOCS-108	61.2	Resistant	HRD	<i>GL BRCA1</i>	-0.1	-	2.3	2.3	2.1	8.2	4	
AOCS-139	62.4	Sensitive	HRD	<i>GL BRCA1</i>	0.11	-	1.6	1.3	2.0	8.0	4	
AOCS-091	39.2	Sensitive	HRD	<i>BRCA1 meth</i>	0.16	-	1.6	0.8	1.1	7.5	4	
AOCS-090	53.8	Sensitive	WT		0.09	-	2.1	1.1	2.2	7.0	4	
AOCS-171	52.1	Sensitive	HRD	<i>SM BRCA1</i>	0.04	-	1.8	1.5	0.7	6.5	4	
AOCS-149	56.6	Sensitive	HRD	<i>SM BRCA2</i>	0.23	-	2.4	1.8	1.6	6.4	4	
AOCS-145	62.2	Sensitive	HRD	<i>GL BRCA1</i>	0.15	-	1.3	1.1	1.5	6.0	4	
AOCS-122	59.0	Sensitive	HRD	<i>SM BRCA2</i>	-0.6	-	0.9	1.0	1.8	6.0	4	
AOCS-116	63.6	Resistant	WT		Gain	0.44	-	2.6	1.4	3.5	5.6	4
AOCS-112	70.0	Resistant	WT		0.26	-	1.2	1.1	3.3	5.5	4	
AOCS-058	58.4	Resistant	HRD	<i>GL BRCA1</i>	0.07	-	0.6	1.2	1.8	5.5	4	
AOCS-125	54.6	Sensitive	HRD	<i>BRCA1 meth</i>	0.16	-	1.2	0.6	1.2	5.4	4	
AOCS-095	50.6	Sensitive	HRD	<i>GL BRCA1</i>	-0.1	-	1.0	0.9	1.4	5.4	4	
AOCS-093	57.7	Sensitive	HRD	<i>BRCA1 meth</i>	0.04	-	1.3	1.4	0.9	5.2	4	
AOCS-147	55.7	Sensitive	HRD	<i>RAD51C meth; SM PTEN del</i>	0.07	-	1.2	6.6	1.1	5.0	2	
AOCS-152	73.0	Sensitive	HRD	<i>SM BRCA1</i>	0.04	-	0.9	1.5	1.5	5.0	4	
AOCS-088	56.5	Sensitive	HRD	<i>GL BRCA1</i>	-0.6	-	0.7	0.7	1.1	4.8	4	
AOCS-170	72.8	Sensitive	HRD	<i>GL CHEK2</i>	0.17	-	2.3	1.6	1.9	4.6	4	
AOCS-158	52.5	Refractory	HRD	<i>BRCA1 meth</i>	0.02	-	0.7	0.9	0.8	4.5	4	
AOCS-114	55.3	Resistant	HRD	<i>GL BRCA1</i>	-0.2	-	0.6	0.8	1.3	4.4	4	
AOCS-128	45.1	Resistant	HRD	<i>BRCA1 meth</i>	0.02	-	0.7	1.9	0.8	4.4	4	
AOCS-107	47.9	Resistant	HRD	<i>BRCA1 meth</i>	0.01	-	1.4	0.8	0.6	4.4	4	
AOCS-113	77.4	Resistant	WT		-0.1	-	1.5	1.2	4.1	4.3	4	
AOCS-105	73.9	Resistant	HRD	<i>GL BRCA1</i>	0.05	-	0.9	0.7	1.4	4.3	4	
AOCS-065	46.3	Sensitive	HRD	<i>GL BRCA1</i>	0.01	-	0.8	0.4	1.0	4.3	4	
AOCS-146	52.1	Sensitive	HRD	<i>GL BRCA1</i>	0.10	-	0.8	1.0	0.9	4.3	4	
AOCS-153	57.9	Refractory	WT		0.03	-	0.9	0.8	1.2	4.2	4	
AOCS-056	72.5	Resistant	WT		0.03	-	0.7	0.6	1.6	4.1	4	
AOCS-131	67.1	Sensitive	HRD	<i>GL BRCA1</i>	0.14	-	0.4	3.1	1.0	4.0	4	
AOCS-034	51.8	Sensitive	HRD	<i>GL BRCA1</i>	0.27	-	1.1	1.4	1.0	3.9	4	
AOCS-148	45.7	Sensitive	HRD	<i>SM PTEN del</i>	0.10	-	0.7	0.9	1.4	3.8	4	
AOCS-092	68.7	Sensitive	HRD	<i>SM BRCA1 del</i>	0.20	-	1.5	0.8	2.8	3.5	4	
AOCS-086	64.7	Sensitive	HRD	<i>SM BRCA1</i>	0.01	-	0.9	0.9	1.5	3.5	4	
AOCS-126	58.9	Sensitive	HRD	<i>GL RAD51C</i>	-0.0	-	0.7	0.9	0.8	3.4	4	
AOCS-104	48.5	Resistant	WT		0.04	-	0.8	0.9	0.9	3.4	4	
AOCS-063	62.3	Resistant	HRD	<i>SM BRCA2</i>	-0.0	-	0.8	0.9	1.4	3.4	4	
AOCS-079	72.8	Resistant	HRD	<i>SM BRCA1</i>	0.02	-	0.7	0.9	1.8	3.4	4	
AOCS-164	55.1	Resistant	HRD	<i>BRCA1 meth</i>	Gain	0.34	-	0.8	1.3	0.9	3.3	4
AOCS-057	56.1	Resistant	HRD	<i>BRCA1 meth</i>	0.13	-	1.0	0.6	1.1	3.3	4	
AOCS-096	61.6	Resistant	WT		0.14	-	1.6	0.5	2.2	3.3	4	
AOCS-004	53.8	Refractory	WT		Gain	0.57	-	1.7	0.6	1.3	3.0	4
AOCS-078	59.9	Resistant	WT		Gain	0.43	-	1.6	2.1	2.1	3.0	4
AOCS-130	65.6	Sensitive	HRD	<i>SM BRCA1</i>	0.07	-	0.7	1.7	1.6	2.9	4	
AOCS-001	53.8	Refractory	HRD	<i>BRCA1 meth</i>	-0.0	-	0.7	1.4	1.0	2.9	4	
AOCS-144	54.4	Sensitive	HRD	<i>BRCA1 meth</i>	0.08	-	0.4	1.3	1.2	2.8	4	
AOCS-160	66.1	Resistant	WT		0.06	-	0.5	0.6	2.8	2.6	3	
AOCS-143	49.7	Sensitive	HRD	<i>GL BRCA1; RAD51C meth</i>	0.16	-	0.5	0.2	1.1	2.5	4	
AOCS-083	58.4	Resistant	WT		0.25	-	1.3	0.7	2.2	2.4	4	
AOCS-165	74.0	Resistant	WT		0.08	-	0.4	0.5	2.9	2.4	3	
AOCS-080	67.6	Refractory	WT		Gain	0.36	-	0.8	0.7	3.1	2.3	3
AOCS-168	59.8	Refractory	WT		-0.1	-	0.6	1.6	1.8	2.2	4	
AOCS-081	59.8	Resistant	WT		Gain	0.54	-	1.3	0.6	2.5	2.2	3
AOCS-076	58.1	Resistant	WT		0.08	-	1.0	0.6	2.2	2.2	3	
AOCS-115	56.4	Resistant	HRD	<i>SM PTEN del</i>	Amplified	0.78	-	0.6	0.3	1.9	2.0	4
AOCS-161	60.6	Refractory	WT		-0.0	-	0.8	1.2	2.5	2.0	3	
AOCS-163	65.6	Resistant	WT		0.31	-	0.5	0.6	2.3	1.8	3	
AOCS-061	59.3	Resistant	WT		Amplified	0.77	-	0.7	0.5	3.0	1.8	3
AOCS-111	65.8	Resistant	WT		Gain	0.46	-	0.8	0.2	2.7	1.5	3
AOCS-084	60.9	Refractory	WT		0.15	-	1.0	0.6	2.0	1.5	3	
AOCS-159	78.4	Resistant	WT		0.01	-	0.7	0.6	2.1	1.3	3	
AOCS-169	74.1	Refractory	WT		0.08	-	0.5	0.4	2.3	1.3	3	
AOCS-005	51.6	Resistant	WT		Amplified	0.90	-	0.9	1.6	1.9	1.2	3
AOCS-133	66.8	Sensitive	HRD	<i>GL BRIP1</i>	0.31	-	0.6	0.4	4.1	1.1	3	
AOCS-123	55.8	Sensitive	WT		-0.0	-	0.5	0.2	1.8	1.1	3	
AOCS-124	58.1	Sensitive	WT		Amplified	0.69	-	0.7	0.4	2.7	1.0	3
AOCS-075	54.0	Resistant	WT		0.27	-	0.4	2.3	1.2	1.0	2	
AOCS-162	74.7	Refractory	WT		Gain	0.37	-	0.7	0.5	2.7	1.0	3
AOCS-109	63.9	Resistant	WT		Amplified	0.80	-	0.7	0.5	3.3	0.8	3
AOCS-157	45.1	Resistant	WT		0.20	-	0.3	0.2	1.6	0.8	3	
AOCS-166	58.0	Refractory	WT	<i>MMR mutation</i>	Amplified	0.67	-	34.4	1.1	1.5	0.7	1
AOCS-085	44.5	Resistant	WT		0.05	-	0.4	0.2	1.0	0.6	3	
AOCS-132	48.7	Sensitive	HRD	<i>GL BRIP1</i>	0.10	-	0.3	0.2	1.7	0.4	3	
AOCS-059	76.7	Resistant	WT		0.00	-	0.6	0.4	2.3	0.4	3	
AOCS-055	59.6	Resistant	WT		Gain	0.55	-	1.2	1.3	0.7	0.4	2
AOCS-097	55.5	Resistant	WT		0.07	-	0.2	0.1	1.1	0.3	3	
AOCS-077	66.5	Refractory	WT		Gain	0.50	-	0.3	0.3	2.1	0.3	3
AOCS-060	69.7	Resistant	WT		Gain	0.49	-	0.3	0.6	3.4	0.1	3

Extended Data Figure 5 | Molecular characteristics of primary tumours. Clinical and molecular characteristics of primary tumours sorted by BRCA signature contribution ($n = 80$).



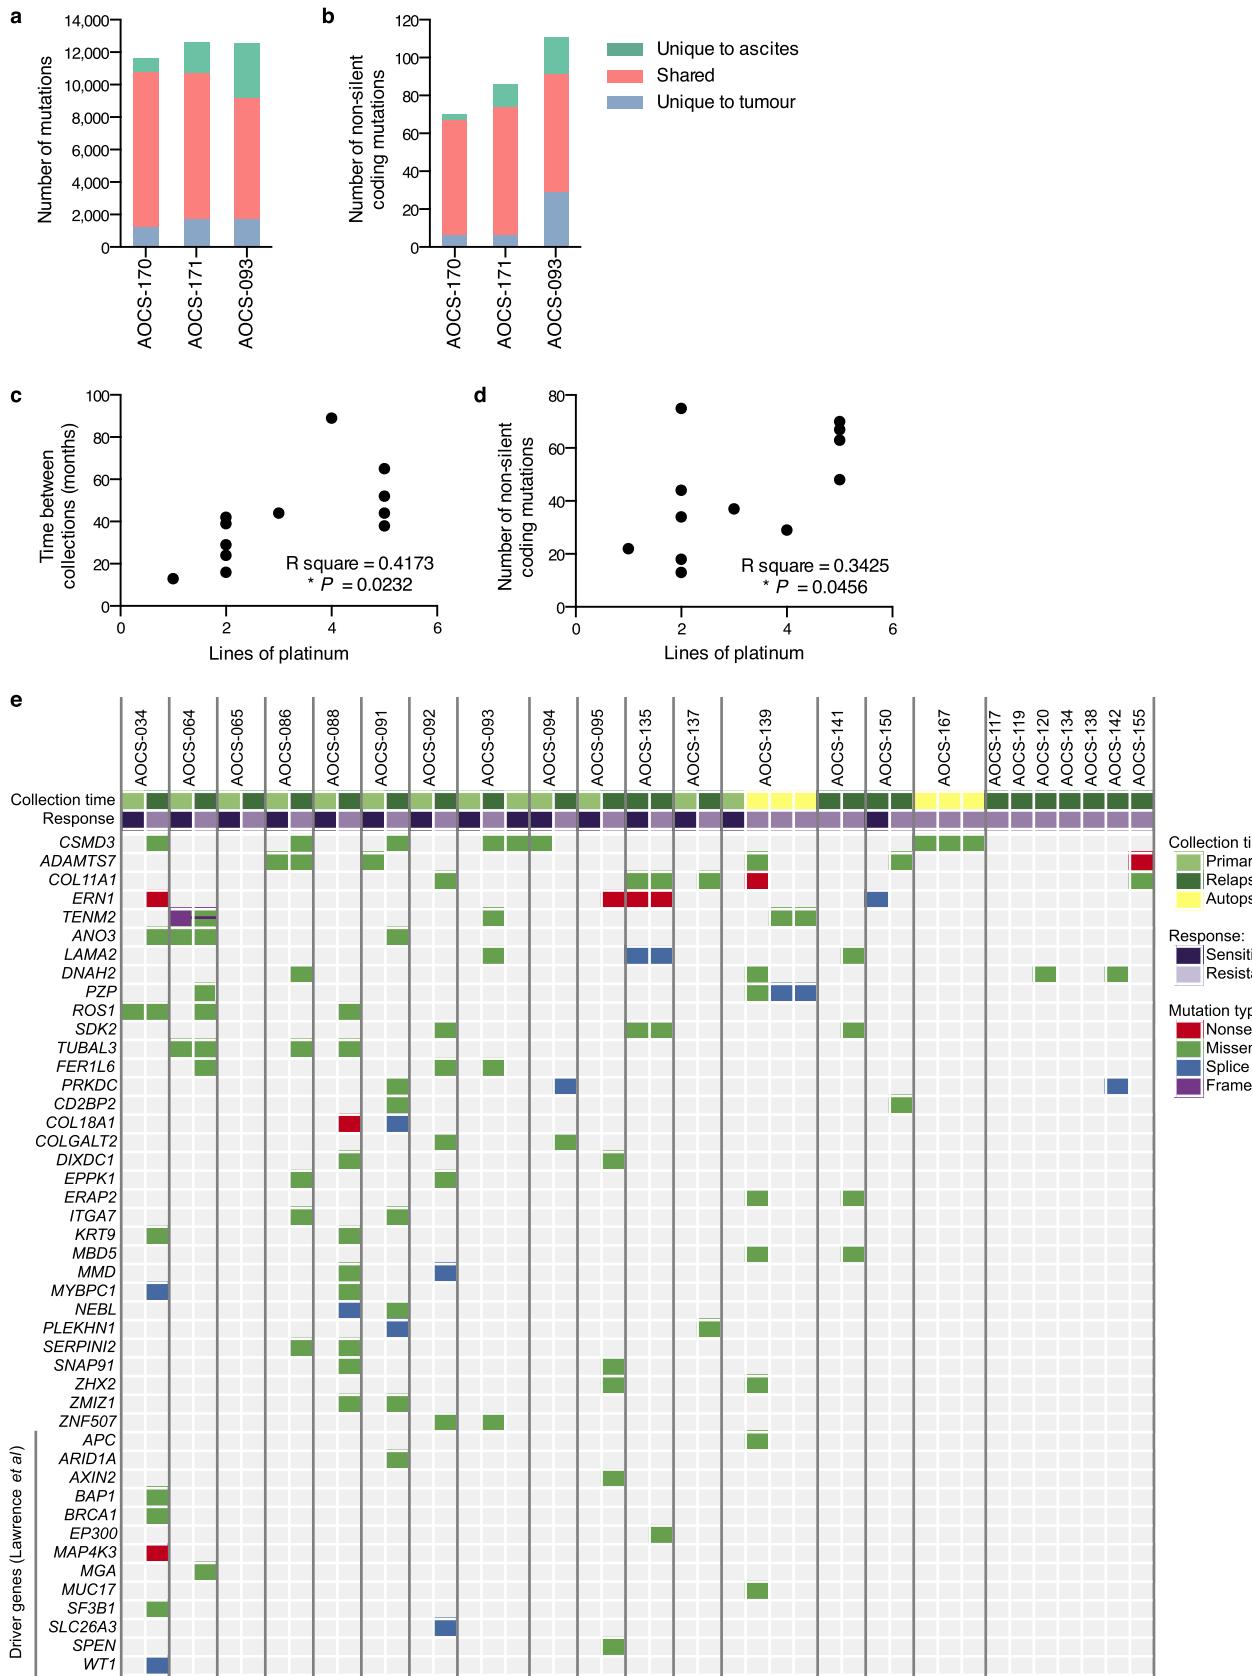
Extended Data Figure 6 | Mutational signature associations. **a, b**, Age mutational signature association with age at diagnosis (**a**) and *CCNE1* copy number (**b**) (Spearman correlation P value reported, $n = 80$). **c, d**, BRCA mutational signature association with *BRCA1/2* and other homologous recombination pathway mutations (**c**) and primary treatment response (**d**) (lines indicate mean; Kolmogorov–Smirnov ** $P < 0.01$, *** $P < 0.001$, **** $P < 0.0001$; $n = 80$).



Extended Data Figure 7 | Molecular drivers and clinical associations.

a, Percentage (*n*) of primary tumours (total *n* = 80) affected by homologous recombination pathway mutations and *CCNE1* copy number gains. One driver mutation counted for samples with more than one change, ranking mutations in *BRCA1/2*, followed by other germline, somatic, amplification, deletion and methylation events respectively. **b**, Association of driver mutation subgroup with overall survival in AOCS and TCGA cohorts (Kaplan–Meier analysis, *P* value calculated by Mantel–Cox log-rank test). **c, d**, Whole genome (c) and

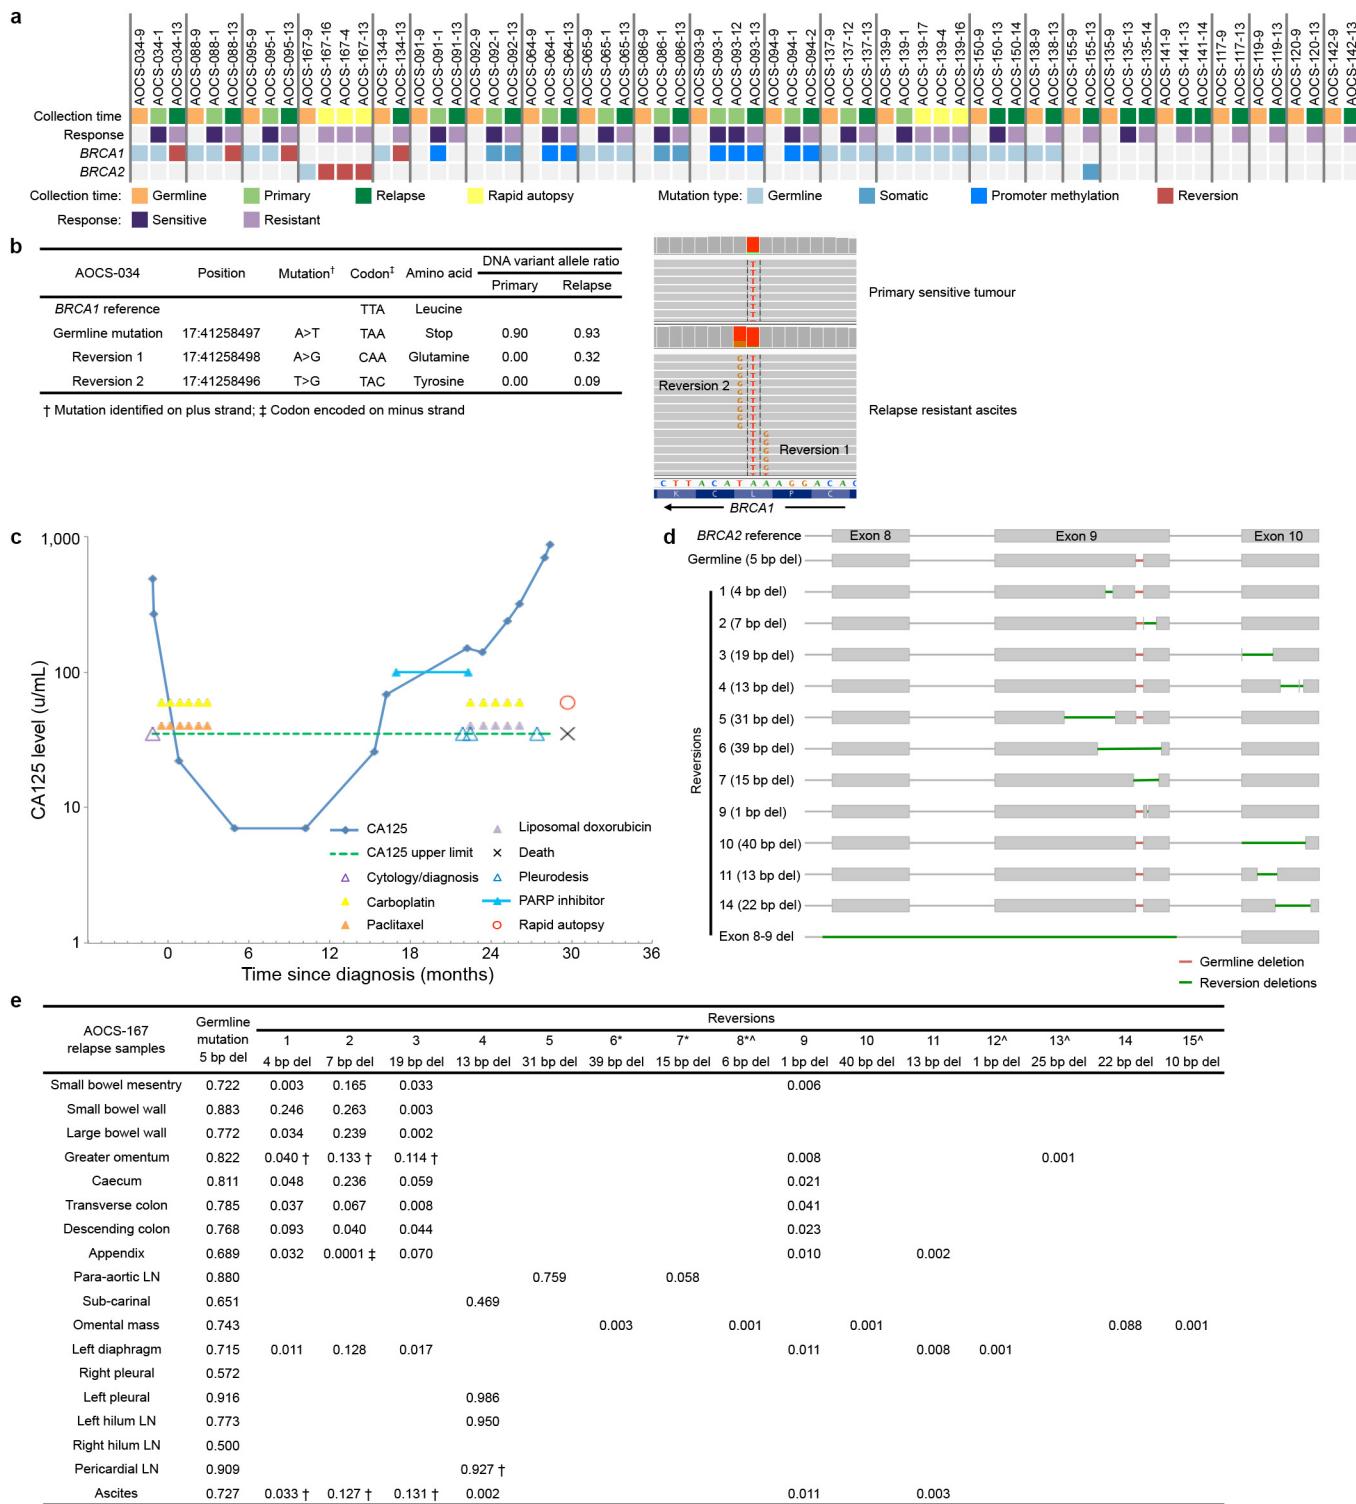
coding mutations (d) per megabase for samples stratified by primary clinical response group (lines indicate mean). Kruskal–Wallis test *P* value reported (**P* < 0.05, ***P* < 0.01, Kolmogorov–Smirnov test). **e**, Boxplots summarize *CCNE1* expression in different driver mutation subgroups (*****P* < 0.0001, unpaired two-tailed *t*-test). Middle bar, median; whiskers, data range. **f**, Proportions of gene expression molecular subtypes between driver mutation subgroups. HR/CCNE1 – subgroup has no detected homologous recombination pathway mutations or *CCNE1* copy number changes.



Extended Data Figure 8 | Analysis of acquired resistance cases.

a, b, Matched primary ascites share most variants with primary tumour samples across the whole genome (**a**) and for non-silent coding mutations (**b**). **c, d**, Significant correlations are observed between the time between collection of the primary and relapse samples and the number of lines of platinum treatment the patient received (Pearson correlation, $P = 0.0232$,

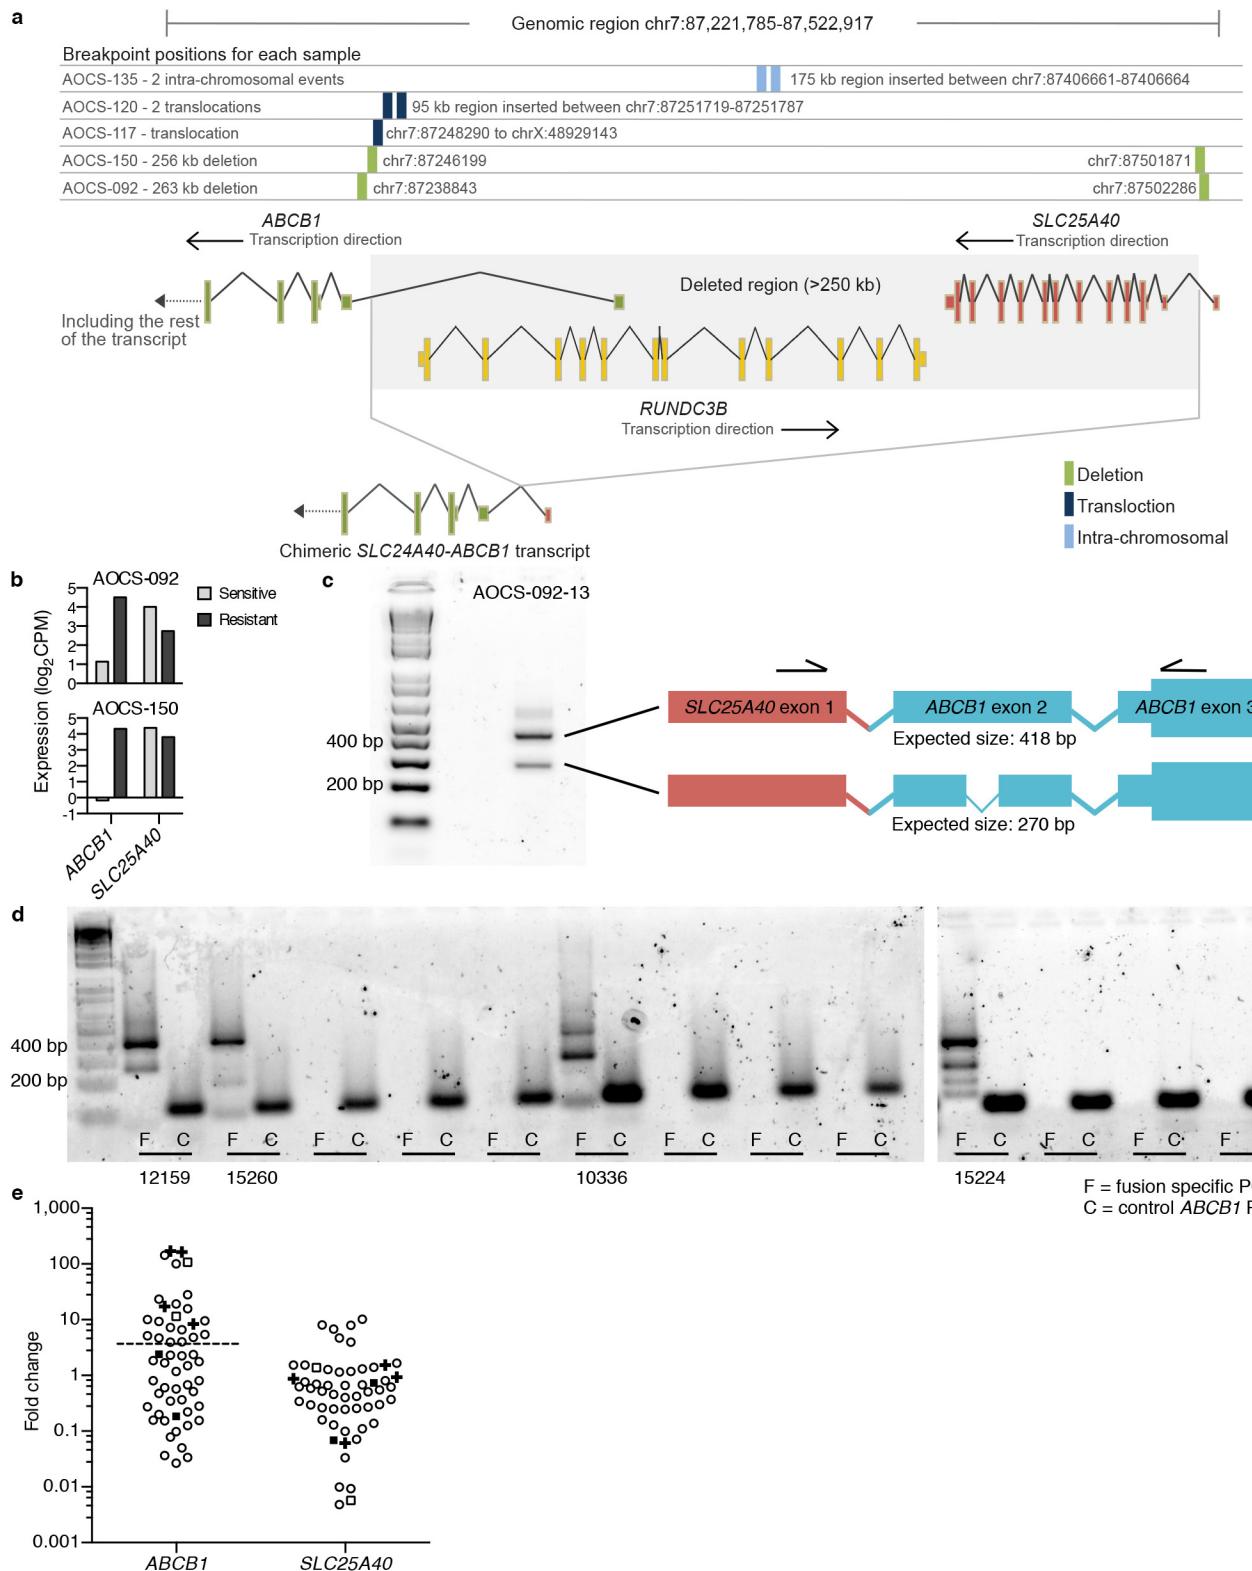
two-tailed) (**c**), and between the number of non-silent coding mutations unique to the relapse samples and the number of lines of platinum treatment the patient received (Pearson correlation, $P = 0.0456$, two-tailed) (**d**). **e**, Genes with recurrent non-silent coding mutations that are unique to the second sample in the acquired resistance cohort or that are driver genes reported previously⁵⁵.



Extended Data Figure 9 | BRCA1/2 reversion events. **a**, *BRCA1* and *BRCA2* mutations are indicated by sample for acquired resistance cases ($n = 23$). Of the ten cases with germline mutations, five show secondary somatic mutations in the relapse-resistant samples that are proposed to restore gene function (that is, reversion). **b**, Two independent point mutations revert the *BRCA1* nonsense mutation in AOCS-034. **c**, CA125 serum marker profile for case AOCS-167.

d, Schematic of 12 high confidence reversion events identified in AOCS-167.

e, Allele frequency of fifteen *BRCA2* reversion events identified in deep amplicon sequencing across 18 relapse samples from case AOCS-167. Reversions observed in single deposits at low allele frequencies (≤ 0.001) were considered low confidence events.



Extended Data Figure 10 | SLC25A40-ABCB1 fusion transcript.

a, Schematic showing the locations of structural variants identified upstream or internal to ABCB1 from WGS. **b**, ABCB1 and SLC25A40 expression levels in sensitive and resistant samples of AOCS-092 and AOCS-150. **c**, RT-PCR verification of SLC25A40-ABCB1 fusion transcript expression in AOCS-092 and schematic of expected RT-PCR products. **d**, RT-PCR results from a validation cohort of relapse ascites identifying four additional samples with the

fusion transcript ($n = 51$). **e**, ABCB1 and SLC25A40 expression levels compared to the SKOV3 cell line by qRT-PCR in the validation cohort and selected cases. Relapse samples with the fusion are indicated by open squares and the matched primary samples indicated by closed squares; validation cases with the fusion are indicated by '+'; the median plus median absolute deviation expression level is indicated by the dotted line ($n = 56$).

A kiloparsec-scale internal shock collision in the jet of a nearby radio galaxy

Eileen T. Meyer^{1,2}, Markos Georganopoulos^{2,3}, William B. Sparks¹, Eric Perlman⁴, Roeland P. van der Marel¹, Jay Anderson¹, Sangmo Tony Sohn⁵, John Biretta¹, Colin Norman^{1,5} & Marco Chiaberge^{1,5,6}

Jets of highly energized plasma with relativistic velocities are associated with black holes ranging in mass from a few times that of the Sun to the billion-solar-mass black holes at the centres of galaxies¹. A popular but unconfirmed hypothesis to explain how the plasma is energized is the ‘internal shock model’, in which the relativistic flow is unsteady². Faster components in the jet catch up to and collide with slower ones, leading to internal shocks that accelerate particles and generate magnetic fields³. This mechanism can explain the variable, high-energy emission from a diverse set of objects^{4–7}, with the best indirect evidence being the unseen fast relativistic flow inferred to energize slower components in X-ray binary jets^{8,9}. Mapping of the kinematic profiles in resolved jets has revealed precessing and helical patterns in X-ray binaries^{10,11}, apparent superluminal motions^{12,13}, and the ejection of knots (bright components) from standing shocks in the jets of active galaxies^{14,15}. Observations revealing the structure and evolution of an internal shock in action have, however, remained elusive, hindering measurement of the physical parameters and ultimate efficiency of the mechanism. Here we report observations of a collision between two knots in the jet of nearby radio galaxy 3C 264. A bright knot with an apparent speed of $(7.0 \pm 0.8)c$, where c is the speed of light in a vacuum, is in the incipient stages of a collision with a slower-moving knot of speed $(1.8 \pm 0.5)c$ just downstream, resulting in brightening of both knots—as seen in the most recent epoch of imaging.

We obtained deep V-band imaging of 3C 264, a radio galaxy 91 megaparsecs (Mpc) from Earth with the Advanced Camera for Surveys (ACS) on board the Hubble Space Telescope (HST) in May 2014. The comparatively deep ACS imaging provides a reference image to compare against previous HST imaging for evidence of proper motions of the four previously known optical knots within the 2"-long jet^{16–18}. We localized over a hundred globular clusters in the host galaxy as a reference system on which to register previous V-band images taken with HST’s Wide Field Planetary Camera 2 (WFPC2) in 1994, 1996, and 2002. The systematic error in the registration of the WFPC2 images is generally of the order of 5 milliarcseconds (mas) or less. After aligning all images to a common reference frame, the fast proper motion of knot B is clearly visible, as shown in Fig. 1 and Supplementary Videos 1 and 2. Previous radio observations have revealed that the initially narrowly collimated jet bends by about 10° at the location marked by the yellow cross¹⁹, which appears to align well with the central axis of the jet in our imaging, and which serves as our reference point for all measured positions.

To measure their apparent speeds, the position of each knot was measured with a centroiding technique. In the case of knots B and C, we also modelled the jet as a constant-density conical jet with superimposed resolved knots, in order to measure their fluxes and positions better, particularly in the final epoch when they appear to overlap. We plotted the position of each knot along the direction of the jet axis

versus time, and fitted the data with a least-squares linear model. A slope significantly larger than zero ($P < 0.05$) indicates significant proper motions, and we used the conversion factor 1.442c years per milliarcsecond to convert angular apparent speeds μ_{app} to dimensionless apparent speed relative to c , β_{app} (see Table 1).

We found that knots A and D have an apparent speed β_{app} consistent with zero (Extended Data Figs 1 and 2), while the inner knots B and C have $\beta_{\text{app}} = 7.0 \pm 0.8$ and 1.8 ± 0.5 , respectively (Fig. 2). The value for knot B exceeds the fastest speeds measured in the jet in the nearby radio galaxy M87, the only other source for which speeds on kiloparsec (kpc) scales have been measured^{20,21}. The current difference in speeds between knots B and C puts them on a collision course, an interaction



Figure 1 | A comparison of HST images of the jet in 3C 264 from 1994 to 2014. The galaxy and core emission from 3C 264 have been subtracted, with a white star representing the location of the black hole, and a yellow cross indicating the location of a bend in the jet seen with radio interferometry¹⁹. The contours show the 30% flux-over-background isophotes around each peak and are overlaid along with vertical guidelines to aid the eye. The first three images were taken with WFPC2, and the final epoch is a deep ACS/WFC image taken for the purposes of measuring proper motions in the jet.

¹Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, Maryland 21218, USA. ²University of Maryland Baltimore County, 1000 Hilltop Circle, Baltimore, Maryland 21250, USA. ³NASA Goddard Space Flight Center, 8800 Greenbelt Road, Greenbelt, Maryland 20771, USA. ⁴Florida Institute of Technology, 150 West University Boulevard, Melbourne, Florida 32901, USA. ⁵Johns Hopkins University, 3400 North Charles Street, Baltimore, Maryland 21218, USA. ⁶Istituto Nazionale Astrofisica, Istituto di Radio Astronomia, Via Piero Gobetti 101, I-40129 Bologna, Italy.

Table 1 | Summary of knot positions and speeds

	Distance along the jet		Proper motion	
	Initial (1994) (mas)	Final (2014) (mas)	μ_{app} (mas yr $^{-1}$)	β_{app}
Knot A	148 ± 5	149 ± 1	0.07 ± 0.20	0.1 ± 0.3
Knot B	257 ± 5	359 ± 2	4.85 ± 0.58	7.0 ± 0.8
Knot C	450 ± 6	478 ± 2	1.27 ± 0.32	1.8 ± 0.5
Knot D	582 ± 5	581 ± 1	-0.13 ± 0.31	-0.2 ± 0.5

Knots A, D by contour method; Knots B, C from maximum-likelihood model.

The distance along the jet is measured from the bend in the radio jet noted in Fig. 1. $\beta_{app} = v/c$, where v is velocity.

which has already begun in the final epoch from 2014, in which the knots begin to overlap (Fig. 2).

In the internal shock model, the collision of two components results in particle acceleration, which will manifest in a significant brightening of the components as they combine into a single moving component. Our modelling results show that in the final 2014 epoch, both knots B and C brighten at the same time by approximately 40% over the mean flux level of the previous three epochs (Fig. 3). The brightening can also be seen in Extended Data Fig. 3, where we show the flux contour along the length of the jet (averaged transversely over a distance of 0.1'') for each epoch.

The flux increase is also corroborated by measuring the flux within a circular aperture centred on and enveloping knots B and C in all epochs, which shows a significant total flux change of approximately 20%–40% over prior epochs. Under equipartition, the cooling length for the optical-emitting electrons for knots B and C is longer than the distance they travelled in our observations, consistent with the lack of any decay in flux levels for these knots over the first three epochs. This is not the case for stationary knot A, which can be seen to decay with a timescale of around 70 years. Knot A appears to be analogous to knot HST-1 in the M87 jet; the latter is thought to be a stationary reconnection shock where the jet pressure drops below that of the external environment²². The event that energized knot A may have been the passage of fast-moving knot B in about 1971^{+8 -17}, yielding a timescale comparable to the knot A decay time.

Previous radio observations of the 3C 264 jet show that the locations of the knots in contemporaneous radio and optical imaging are very similar^{18,23}. The earliest radio image of 3C 264 taken in October 1983 with the VLA clearly shows three distinct features (see contour comparison in Extended Data Fig. 4). While knot D appears to be completely stationary over this time-frame, the data suggest the possibility

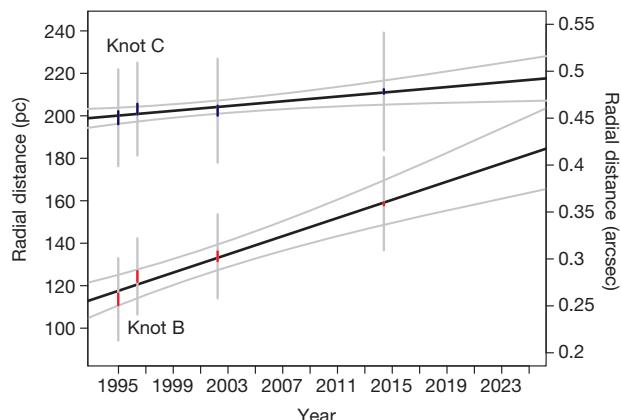


Figure 2 | Position versus time for knots B and C. The position of the centre of each knot is noted by thick red and blue lines, respectively, with line extent corresponding to the $3\sigma \chi^2$ modelling error plus the systematic error on the mean from the image registration. The grey extensions show the best-fit size of the resolved knot from the model. As shown, the best-fit linear slope yields a speed of 7 ± 0.8 for knot B and 1.8 ± 0.5 for knot C. In the final epoch the knots are directly adjacent and possibly overlapping, within the errors.

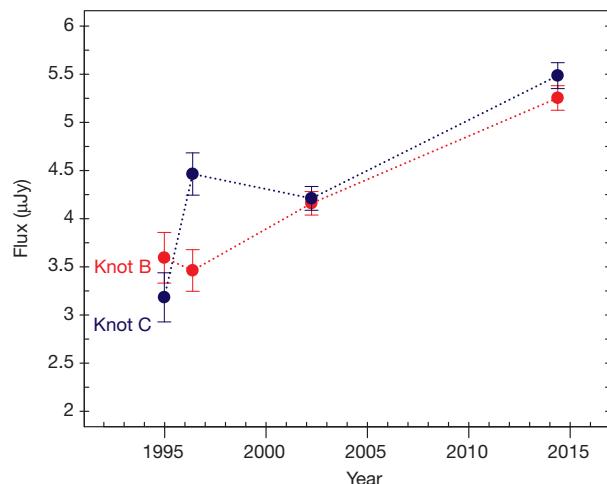


Figure 3 | Change in optical flux at 6,000 Å in the colliding knots B and C over 20 years. The knots show a simultaneous significant increase in flux in the 2014 image. Errors are the $3\sigma \chi^2$ modelling error.

that knots B and C were moving faster in the past and may have decelerated (Extended Data Fig. 5).

In the internal shock model, the efficiency η of the conversion of the dissipated kinetic energy E_{diss} into radiation is generally unknown. For the collision in 3C 264 we can estimate it from our observations (further details are in the Methods): we take the minimum possible value of the Lorentz factor $\Gamma_B = 7.1$ for knot B so that the jet is observed at an angle of 8.1° , and for slower knot C we take $\Gamma_C = 2.8$. We assume an equipartition magnetic field and one proton per electron in the plasma to obtain masses of 2.5×10^{30} g and 6.7×10^{30} g for knots B and C. After the collision, the single combined component will move with Lorentz factor $\Gamma_m = 3.7$. The total energy dissipated is taken as the difference in kinetic energy before and after the collision and is $E_{diss} = 2.6 \times 10^{51}$ erg. If the knot B+C complex stays at the flux level observed in 2014 throughout the duration of the collision (which we take to be the knot superposition time of 30 years in the observer's frame, as one would expect if the cooling time of the optically emitting electrons is shorter than the collision time), the efficiency of conversion is only $\eta = 10^{-3}$, considerably lower than is usually assumed in models.

This, however, is a lower limit for two reasons. First, theoretical modelling²⁴ of internal shocks suggests that the flux should steadily rise to a peak occurring halfway through the collision, which would increase η . Second, it is possible that the cooling time of the optically emitting electrons is longer than the collision time, which would increase the duration of the elevated optical-ultraviolet emission, and therefore η . The rate of cooling depends on the magnetic field in the shocked plasma. Long-term monitoring of the collision in 3C 264 over the coming decades could probe the evolution of the flux and hence constrain two free parameters of the internal shock model: the fraction of E_{diss} that goes to radiating electrons and the fraction that goes to generating the magnetic field in the shocked plasma.

Online Content Methods, along with any additional Extended Data display items and Source Data, are available in the online version of the paper; references unique to these sections appear only in the online paper.

Received 25 November 2014; accepted 31 March 2015.

1. Livio, M. Astrophysical jets: a phenomenological examination of acceleration and collimation. *Phys. Rep.* **311**, 225–245 (1999).
2. Rees, M. J. The M87 jet—internal shocks in a plasma beam. *Mon. Not. R. Astron. Soc.* **184**, 61–65 (1978).
3. Medvedev, M. V. & Loeb, A. Generation of magnetic fields in the relativistic shock of gamma-ray burst sources. *Astrophys. J.* **526**, 697–706 (1999).
4. Paczynski, B. & Xu, G. Neutrino bursts from gamma-ray bursts. *Astrophys. J.* **427**, 708–713 (1994).
5. Rees, M. J. & Meszaros, P. Unsteady outflow models for cosmological gamma-ray bursts. *Astrophys. J.* **430**, L93–L96 (1994).

6. Fender, R. P., Belloni, T. M. & Gallo, E. Towards a unified model for black hole X-ray binary jets. *Mon. Not. R. Astron. Soc.* **355**, 1105–1118 (2004).
7. Spada, M., Ghisellini, G., Lazzati, D. & Celotti, A. Internal shocks in the jets of radio-loud quasars. *Mon. Not. R. Astron. Soc.* **325**, 1559–1570 (2001).
8. Fomalont, E. B., Geldzahler, B. J. & Bradshaw, C. F. Scorpius X-1: the evolution and nature of the twin compact radio lobes. *Astrophys. J.* **558**, 283–301 (2001).
9. Migliari, S., Fender, R. & Méndez, M. Iron emission lines from extended X-ray jets in SS 433: reheating of atomic nuclei. *Science* **297**, 1673–1676 (2002).
10. Abell, G. O. & Margon, B. A kinematic model for SS433. *Nature* **279**, 701–703 (1979).
11. Hjellming, R. M. & Rupen, M. P. Episodic ejection of relativistic jets by the X-ray transient GRO J1655 – 40. *Nature* **375**, 464–468 (1995).
12. Mirabel, I. F. & Rodríguez, L. F. A superluminal source in the Galaxy. *Nature* **371**, 46–48 (1994).
13. Dhawan, V., Mirabel, I. F. & Rodriguez, L. F. AU-scale synchrotron jets and superluminal ejecta in GRS 1915+105. *Astrophys. J.* **543**, 373–385 (2000).
14. Cheung, C. C., Harris, D. E. & Stawarz, Ł. Superluminal radio features in the M87 jet and the site of flaring TeV gamma-ray emission. *Astrophys. J.* **663**, L65–L68 (2007).
15. Marscher, A. P. *et al.* Probing the inner jet of the quasar PKS 1510–089 with multi-waveband monitoring during strong gamma-ray activity. *Astrophys. J.* **710**, L126–L131 (2010).
16. Crane, P. *et al.* Discovery of an optical synchrotron jet in 3C 264. *Astrophys. J.* **402**, L37–L40 (1993).
17. Baum, S. A. *et al.* HST and Merlin observations of 3C 264—a laboratory for jet physics and unified schemes. *Astrophys. J.* **483**, 178–193 (1997).
18. Perlman, E. S. *et al.* A multi-wavelength spectral and polarimetric study of the jet of 3C 264. *Astrophys. J.* **708**, 171–187 (2010).
19. Lara, L., Giovannini, G., Cotton, W. D., Feretti, L. & Venturi, T. The inner kiloparsec of the jet in 3C 264. *Astron. Astrophys.* **415**, 905–913 (2004).
20. Biretta, J. A., Sparks, W. B. & Macchetto, F. Hubble Space Telescope observations of superluminal motion in the M87 jet. *Astrophys. J.* **520**, 621–626 (1999).
21. Meyer, E. T. *et al.* Optical proper motion measurements of the M87 jet: new results from the Hubble Space Telescope. *Astrophys. J.* **774**, L21 (2013).
22. Stawarz, Ł. *et al.* Dynamics and high-energy emission of the flaring HST-1 knot in the M 87 jet. *Mon. Not. R. Astron. Soc.* **370**, 981–992 (2006).
23. Lara, L. *et al.* The radio-optical jet in NGC 3862 from parsec to subkiloparsec scales. *Astrophys. J.* **513**, 197–206 (1999).
24. Kobayashi, S., Piran, T. & Sari, R. Can internal shocks produce the variability in gamma-ray bursts? *Astrophys. J.* **490**, 92–98 (1997).

Supplementary Information is available in the online version of the paper.

Acknowledgements E.T.M. acknowledges HST grant GO-13327.

Author Contributions E.T.M. performed the HST data analysis and wrote the paper. B.S., J.A., R.P.v.d.M. and S.T.S. were consulted on and contributed to the HST data analysis. M.G. contributed to the interpretation and performed the theoretical calculations in consultation with E.T.M. E.P. contributed radio data. J.B., C.N. and M.C. provided insight into the design of the observations and interpretation. All authors discussed the results and commented on the manuscript.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to E.T.M. (eileen.meyer@gmail.com).

METHODS

HST data analysis. These methods are essentially identical to that used in the previous study of M87, and we refer to ref. 21 for more extensive details beyond the outline given here. The imaging used for this study consisted of two orbits of F606W ACS/WFC imaging obtained in 2014 (eight exposures of approximately 600 s each, proposal ID 13327), and the following WFPC2 planetary camera (PC) imaging: nine 140-s exposures in the F702W filter in 2002 (proposal ID 9069), one 350-s and one 400-s exposure taken in the F791 filter in 1996 (proposal ID 5927), and two 140-s exposures in the F702W filter taken in 1994 (proposal ID 5476).

To build the reference frame, in each ACS exposure approximately 100 common globular clusters were detected and localized using a point spread function (PSF) peak-fitting routine. Taking into account the filter-dependent geometric distortion, a linear six-parameter transformation solution was found for each individual ACS exposure based on these positions into a common master frame. The final systematic error in the registration of the ACS exposures is about 0.007 pixels or 0.18 mas. The WFPC2/PC imaging was similarly matched to the master reference frame, using approximately 10–15 bright globular clusters per raw image. The final systematic error in the registration of the WFPC2/PC images is of the order of about 0.14 pixels or 3.5 mas. We include a systematic error of 5 mas for all WFPC2 images and 0.2 mas for the ACS image, added in quadrature to the measured positional errors in each epoch. For each of the four epochs, the transformation solutions were used to stack the exposures into a mean image (25 mas per pixel resolution), correcting for exposure time differences (scaled to the longest exposure), and iteratively sigma-clipping to reject cosmic rays (when more than two images were used). The final image stacks are referred to below by the year of the observations (1994, 1996, 2002, and 2014 images).

The host galaxy of 3C 264 is elliptical, with a light distribution increasing smoothly towards the centre up to the inner 0.9'', where a dust disk is clearly present²⁵ and absorbing some of the galaxy light. The host galaxy light outside the dust disk was modelled individually for the 1996, 2002, and 2014 images using the IRAF (image reduction and analysis facility) software tasks *ellipse* and *bmodel*, and applying a selective mask to avoid the regions of the jet and globular clusters (Extended Data Fig. 6). The 2002 model was also used for the 1994 epoch since the former is deeper. The dust disk was modelled individually for each epoch, by measuring the flux along progressively smaller contour line circles around the centre of the galaxy. The model flux was taken to be a moving average around any individual pixel, interpolating over masked-out areas. The dust disk was modelled up to two pixels (50 mas) from the centre of the galaxy. Inside this radius it is very difficult to reliably separate the core of the jet (a point source) from the galaxy owing to the irregular dust disk, which prevents fitting of the galaxy with a Sérsic or exponential model. Since the core flux level is not of primary interest for this study, we simply take the image values inside this radius as the model, and then subtract the entire model of galaxy and dust disk from our images for each epoch. Finally, we scaled all of the images to flux units (in μJy) at a reference frequency of 6,000 Å using the IRAF task *calphot*, assuming a power-law form ($f_v \propto v^{-1}$), $E(B - V) = 0.02$, and making the correction for sensitivity changes between decontamination events.

Proper motions. For all knots (A–D) and all epochs, we applied the same contour technique to measure the locations of each knot as applied in the M87 study²¹, except that the contour line levels were drawn at 30% flux over background. In the final 2014 epoch, knots B and C were inside a single contour. The knots are also resolved and cannot be easily modelled as point sources to calculate their flux level. To better measure the fluxes of knots and obtain independent locations for the knots, we modelled the jet as a two-dimensional image, focusing only on the region of knots A, B and C. The model, described below, has 14 parameters, and our maximum-likelihood optimization yields an estimate of the knot flux and location relative to the core.

Maximum-likelihood modelling of the 3C 264 jet. In all epochs, the model was optimized to match the two-dimensional image over a subset of the image plane centred on knots B and C (Extended Data Fig. 7), covering a total of 226 reference frame pixels, which we call the ‘image box’. We fixed the central jet axis at a position angle (North through East) of 40.43° , running from one end of the image box to the other. The ‘background’ jet is modelled as a cone parameterized by radius l_1 at the distance from the core coinciding with the start of the image box and increasing linearly to l_2 at the other end of the box. The intensity of the jet (in μJy per pixel) is characterized by a single value I_{jet} . The effective intensity at any given point changes both linearly along the jet (in such a way as to mimic an expanding jet with a constant rate of flow) and laterally from the jet axis, to mimic a decreasing line-of-sight depth through the jet. Given any point inside the defined jet area, we have w_{jet} as the radius of the jet at that distance from the core, and x_{jet} as the

perpendicular distance of the particular point from the axis. The specific intensity is thus

$$I = I_{\text{jet}} \frac{1}{w_{\text{jet}}} \sqrt{(w_{\text{jet}}^2 - x_{\text{jet}}^2) / w_{\text{jet}}^2} \quad (1)$$

On top of this jet, we model knots A, B, and C as circular disks parameterized by radii (r_A , r_B , r_C), distances from the core (d_A , d_B , d_C), and flux densities in μJy per pixel (I_A , I_B , I_C) which is constant over the disk and zero outside. Knots B and C were also allowed to move in the transverse direction perpendicular to the jet axis. Knot D is not included because the image box does not include this area. The full jet model thus has 14 free parameters. The simulated image is created by adding the contribution of each of the four components to each individual pixel, before applying the relevant PSF.

To accurately fit the jet model to the image data, the simulated image must be convolved with an appropriate PSF, which is different for each image epoch. We used the TinyTim²⁸ software package to create model PSFs for both WFPC2/PC and ACS/WFC using the appropriate filter to match the observations. We chose a roughly central pixel location—(400, 400) for WFPC2/PC or (1,000, 3,000) for ACS/WFC—standard focus, and a power-law object spectrum with $F_v \approx v^{-1}$, the spectral index of -1 being that previously observed for the knots in the jet¹⁸. For the WFPC2/PC PSFs, the initial PSF was subsampled by a factor of ten before being rebinned into the WFPC2/PC native PC scale, smoothed with the recommended 3×3 pixel Gaussian, and finally resampled onto the $0.025''$ pixel scale of our stacked images. Optimal values of the 14 free parameters which lead to a simulated image that best fits the data were found by the method of maximum likelihood. We treated each pixel as an independent Poisson trial with a predicted value λ_i and measured value x_i , where

$$\lambda_i = g \frac{p_i}{a_{\text{scale}}} + b_i \quad (2)$$

and

$$x_i = g \frac{f_i}{a_{\text{scale}}} + b_i \quad (3)$$

Here g is the gain ($g = 7$ electrons per DN, where DN is data number, for the WFPC2/PC images, $g = 1$ for the ACS/WFC images) and p_i is the predicted flux in the simulated image pixel. The constant a_{scale} is the scaling factor, specific to each epoch, which was originally used to multiplicatively scale the images from units of DN or electrons to flux in μJy , while b_i is the (unscaled) background estimation that was subtracted from the images before scaling. The same equation is used to calculate x_i except that we use the measured flux in each pixel f_i rather than the simulated prediction.

The maximum-likelihood method requires that we maximize the likelihood function, which is simply the product of the probability density values for each pixel. Because the number of total counts is large (typically $\gg 1,000$) and in order to accurately calculate the likelihood contribution for each pixel, we use the normal approximation to the Poisson probability distribution in calculating the likelihood. We also chose to maximize the log of the likelihood function, which is equivalent to maximizing the likelihood but more practical to implement. Given the above definitions of λ_i and x_i for the i th pixel in the image box, the contribution to the likelihood is simply

$$pr_i = \frac{1}{\sqrt{2\pi\lambda_i/n}} e^{-\frac{(\lambda_i - x_i)^2}{2\lambda_i/n}} \quad (4)$$

where n is the total number of images in the stack. The final log-likelihood function to be maximized is thus

$$\mathcal{L} = \sum_i \log(pr_i)^{1/f} \quad (5)$$

where f is the correction for the different pixel scales between WFPC2/PC and ACS/WFC and our super-sampled reference frame described in the methods, with $f = 3.856$ for WFPC2/PC and $f = 3.9204$ for ACS/WFC.

Although a number of approaches are possible for finding the optimal values of the 14 free parameters that will maximize \mathcal{L} , we chose to use the freely available Simulated Annealing code²⁶, which is well-suited to multi-parameter optimizations of non-smooth functions that are likely to have many false maxima ‘traps’. All 14 parameters were allowed to roam inside a wide range of reasonable values, and we confirmed that the final optimal value did not show signs of running into upper or lower bounds.

The final results of the jet modelling are given in Extended Data Table 1. We have used the data for knots B and C only in the main part of the paper, as the subtraction of the background on the inner side of knot A is somewhat unreliable due to the nearness of the core, leading us to place one edge of the image box over the inner side of knot A. Knot A was thus included in the model more for

improving the fit of the other components than for measurement. In the main text, location values for knots A and D are taken from the contour method, while for knots B and C are taken from the best-fit model (radial positions only). Note that the transverse motions are only marginally significant once the systematic registration error is included, except in the final epoch where both knots are slightly deviated to the North.

Errors for each parameter were calculated by freezing all other parameters to their best-fit values, and progressively increasing (or decreasing) the value of the parameter of interest. This allowed us to plot the change in the likelihood function $\Delta\mathcal{L}$ versus the parameter value. The shape of this function is roughly parabolic with a peak at the optimal parameter value. The distribution of $2\Delta\mathcal{L}$ is a χ^2 variable with one degree of freedom. Thus an upper and lower error value can be calculated by finding the parameter value corresponding to a 3σ value of $2\Delta\mathcal{L}$. The errors reported in Extended Data Table 1 are the 3σ errors. Since the errors are two-sided, the criterion used was that $2\Delta\mathcal{L} = 10.27288$. Fluxes were calculated by multiplying the flux density best-fit value by πr^2 where r is the best-fit radius, and appropriately propagating the errors.

The flux increase implied by the modelling results can be checked using aperture photometry. We used a 6.5-pixel circular aperture, centred on knots B and C so that both are almost entirely contained within the aperture during all four epochs without moving it. Using the 2014 modelling results, we estimate that the smooth-jet contribution to the flux in this aperture is $17.4 \mu\text{Jy}$. The total fluxes attributed to the jet for the four epochs are $7.9 \mu\text{Jy}$, $7.8 \mu\text{Jy}$, $6.7 \mu\text{Jy}$, and $9.4 \mu\text{Jy}$, indicating a 20%–40% flux change in the final epoch. The estimated error on these fluxes is approximately $0.3 \mu\text{Jy}$, where errors from the Poisson noise ($0.2 \mu\text{Jy}$ and $0.05 \mu\text{Jy}$ for WFPC2 and ACS respectively), irregularities in the dust disk ($0.15 \mu\text{Jy}$), and possible contributions of knot A to the aperture ($0.2 \mu\text{Jy}$) have been added in quadrature. Note that these estimates do not correct for the flux which falls outside the aperture, and are thus slightly less than the totals for knots B and C taken from Extended Data Table 1.

The extended baseline with VLA observations. 3C 264 was observed in October 1983 at U band (14.9 GHz) in A configuration as part of program AB239. The quasar 3C 286 was used as the primary flux density calibrator source while QSO 1144+402 was used for a primary phase calibrator and secondary flux density calibrator. The observations used two intermediate frequencies in each band, each with a bandwidth of 50 MHz, and the average of the two was used to produce the final maps. The reduction of the 1983 U-band data followed closely the recommended procedure in chapter 4 of the AIPS Cookbook (<http://www.aips.nrao.edu>). Very little flagging was required. Once calibrated, a clean map was made with several rounds of self-calibration, and deconvolution using the AIPS tasks *clean* and *calib*. We used the ‘robust’ or Briggs’ weighting scheme to weight the visibilities.

The earlier radio image clearly shows three distinct features (see contour comparison in Extended Data Fig. 4). In the case of knot B, a quadratic fit to the HST and VLA data suggests that in late 1983 the knot had a speed of $10.2c$, slowing by $0.16c$ per year to reach $5.6c$ at the beginning of 2014. While a quadratic fit is very poor for the combined data on knot C, a linear fit between the 1983 and 1994 epochs similarly suggests a speed of $9.8c$. Strong conclusions are prevented by the disparity in imaging frequencies and the large errors on the radio positions.

Theoretical considerations. We assume that the flow consists of an electron-proton plasma, with one proton per electron, in energy equipartition with the comoving magnetic field. We derive the equipartition magnetic field values at knots B and C using standard expressions including beaming²⁷. We use the 2002 flux values at 22.46 GHz of 4.4 mJy and 4.1 mJy , and diameters of 39.6 pc and 48.9 pc , respectively, for knots B and C, as well as the observed radio spectral index of $\alpha = 0.5$. We assume that the knots are spherical, noting that the measured diameter in the comoving frame for a spherical blob equals that in the observer’s frame. We assume a low-frequency cutoff of 10 MHz in the observer’s frame. We

derive values of $1.2/\delta \text{ mG}$ and $1.0/\delta \text{ mG}$ (where δ is the Doppler factor) for knots B and C respectively. These values are nearly identical to the values given in ref. 23 once corrected for beaming (they assumed a jet angle to the line of sight $\theta = 50^\circ$, which has now been ruled out by the superluminal motions we observe).

We calculated masses $m_B = 2.5 \times 10^{30} \text{ g}$ and $m_C = 6.7 \times 10^{30} \text{ g}$ for knots B and C respectively, through $m = 2U_B V/c^2$, where U_B is the energy density of the magnetic field in equipartition and V is the volume of the knot in the comoving frame. Using these masses and the Lorentz factors $\Gamma_B = 7.1$, $\Gamma_C = 2.8$ (as derived by the relation $\Gamma_{B,\min} = (\beta_{B,\text{app}}^2 + 1)^{1/2}$, and assuming that $\Gamma_B = \Gamma_{B,\min}$, which sets the jet angle θ to the line of sight to $\arcsin \theta = 1/\Gamma_B = 8.1^\circ$), the Lorentz factor Γ_m of the single component resulting from the collision is²⁷

$$\Gamma_m = \left(\frac{m_B \Gamma_B + m_C \Gamma_C}{m_B/\Gamma_B + m_C/\Gamma_C} \right)^{1/2} = 3.7 \quad (6)$$

as found by conservation of momentum. The energy that will be dissipated cumulatively throughout the collision is

$$E_{\text{diss}} = m_B(\Gamma_B - \Gamma_m) + m_C(\Gamma_C - \Gamma_m) = 2.6 \times 10^{51} \text{ erg} \quad (7)$$

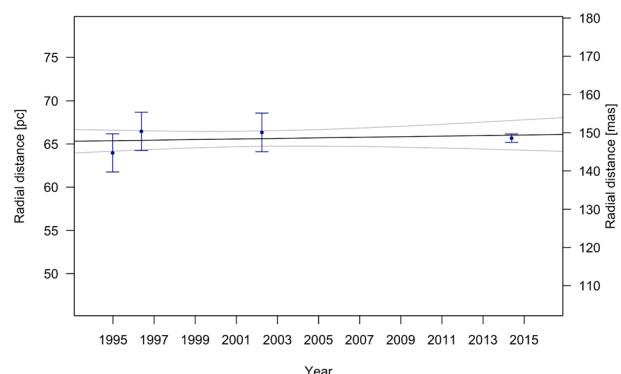
Assuming that the collision lasts from the moment the collision commences up to the point where the back end of the fast component has reached the back end of the slow component, the duration of the collision in terms of coordinate time in the galaxy frame is $\Delta t = 2R'/[\Gamma_B(\beta_B - \beta_C)c] \approx 340$ years, where R' is the comoving radius of the fast component, while the duration for the observer is $\Delta t_{\text{obs}} = 2R' / (\Gamma_B \Gamma_C \delta_C (\beta_B - \beta_C)c) \approx 25$ years, in approximate agreement with the 30-year time required for the knots to superimpose based on the observed HST proper motions. The power dissipated in the galaxy frame is then $P_{\text{diss}} = E_{\text{diss}}/\Delta t = 2.4 \times 10^{41} \text{ erg s}^{-1}$. If the dissipated energy was completely converted to radiation, the solid-angle integrated luminosity in the galaxy frame $L_{4\pi}$ would be equal to P_{diss} . We can estimate the current $L_{4\pi}$ due to the collision by using the flux increase f_{coll} in knots B and C from 2002 to 2014, to obtain the luminosity increase assuming isotropy $L = 1.3 \times 10^{40} \text{ erg s}^{-1}$, which corresponds to a solid-angle integrated luminosity $L_{4\pi} = L I_{\text{m}}^{1/2}/\delta_m^4 = 1.6 \times 10^{38} \text{ erg s}^{-1}$. The efficiency of converting the available energy to radiation is then $\eta = L_{4\pi}/P_{\text{diss}} \approx 10^{-3}$; as explained in the main text, this is a lower limit.

The cooling length of a knot. Using standard formulae, the comoving frame cooling time at an observed frequency v of a knot in equipartition is

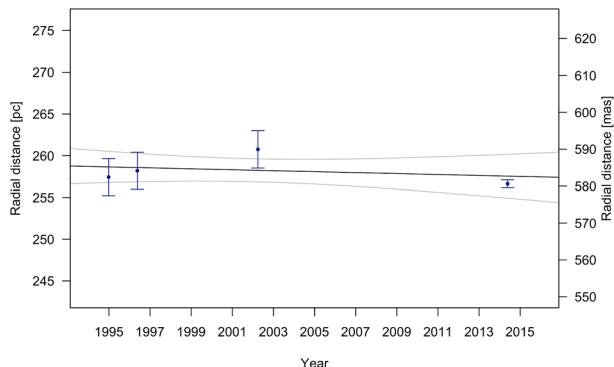
$$t' = \frac{6\pi m_e c}{\sigma_T B_{\text{eq},0}^2} \left(\frac{m_e c^2 B_{\text{eq},0}}{h B_{\text{cr}} v} \right)^{1/2} \delta^2 = t_0 \delta^2 \quad (8)$$

where $B_{\text{eq},0}$ is the equipartition magnetic field assuming $\delta = 1$, m_e is the electron mass, σ_T is the Thomson cross-section, h is Planck’s constant, $B_{\text{cr}} = 2\pi m_e^2 c^3/(eh) = 4.4 \times 10^{13} \text{ G}$ is the critical magnetic field and e is the electron charge. This corresponds to a propagation length of $l = c\beta \Gamma t' = c\beta \Gamma \delta^2 t_0$, which in turn corresponds to a projected cooling length $l_{\text{obs}} = ls\sin\theta = t_0 c \delta \beta_{\text{app}}$, where we have used the standard expression $\beta_{\text{app}} = \beta \sin\theta / (1 - \beta \cos\theta)$. Using values appropriate for knot B we obtain $l_{\text{obs}} = 874 \text{ pc}$ and for knot C $l_{\text{obs}} = 104 \text{ pc}$, in both cases longer than the propagation length recorded by our observations (about 43 pc for knot B and about 11 pc for knot C).

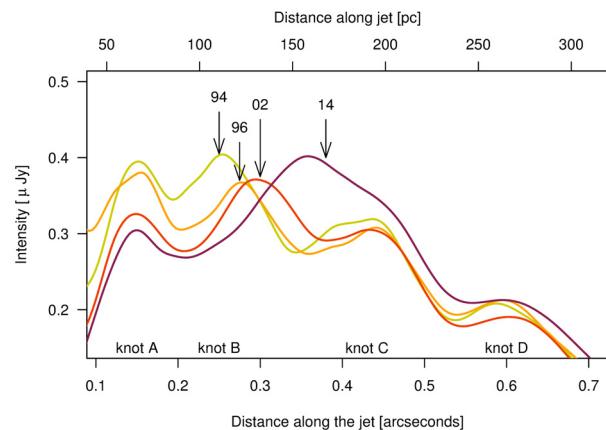
25. Sparks, W. B., Baum, S. A., Biretta, J., Macchetto, F. D. & Martel, A. R. Face-on dust disks in galaxies with optical jets. *Astrophys. J.* **542**, 667–672 (2000).
26. Goffe, W. L., Ferrier, G. D. & Rogers, J. Global optimization of statistical functions with simulated annealing. *J. Econ.* **60**, 65–99 (1994).
27. Harris, D. E. & Krawczynski, H. X-ray emission processes in radio jets. *Astrophys. J.* **565**, 244–255 (2002).
28. Krist, J. E., Hook, R. N. & Stoehr, F. 20 years of Hubble Space Telescope optical modeling using Tiny Tim. *Proc. SPIE* **8127** (2001).



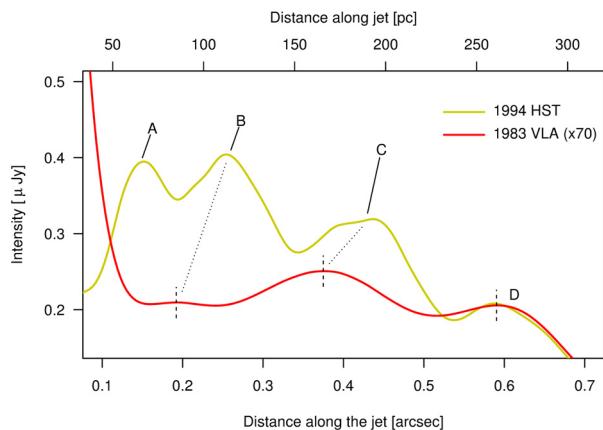
Extended Data Figure 1 | Position versus year for knot A. Positions are taken from the contour analysis method. Errors are 1σ from contour-derived position measurement added to the systematic error on the mean position from the image registration in quadrature.



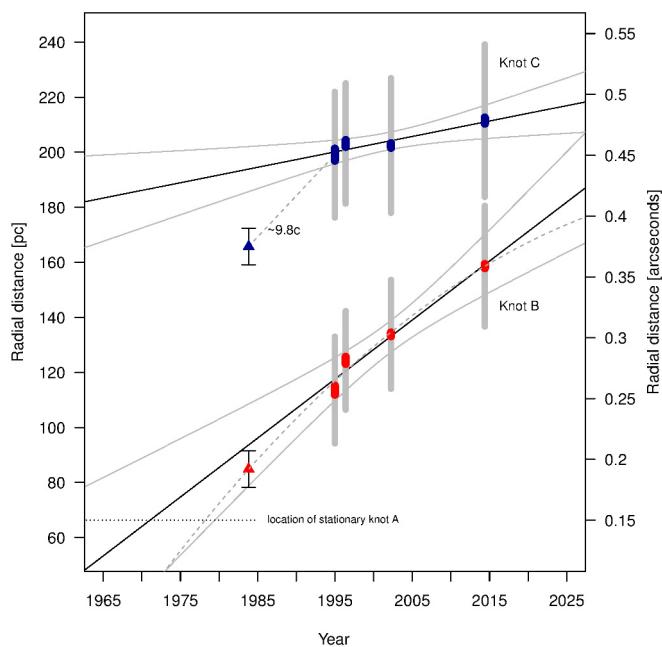
Extended Data Figure 2 | Position versus year for knot D. Positions are taken from the contour analysis method. Errors are 1σ from contour-derived position measurement added to the systematic error on the mean position from the image registration in quadrature.



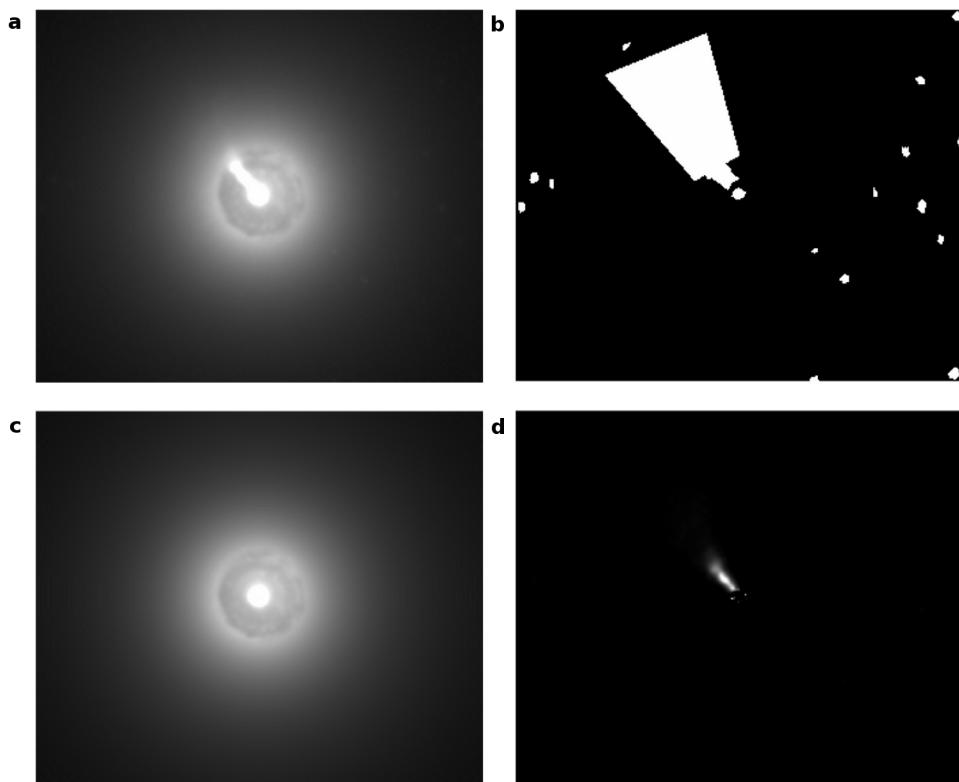
Extended Data Figure 3 | Mean flux along the midline of the jet. Flux is averaged transversely over $0.1''$ at each step along the jet. The epoch is denoted by arrows. Note the decay of knot A with time, while the knot B+C complex appears substantially brighter in the last epoch.



Extended Data Figure 4 | Comparison of 1983 VLA midline flux to 1994 HST midline flux. Flux is averaged transversely over $0.1''$ at each step along the jet. We fit parabolic forms to the peaks seen in the radio contour to derive rough estimates of knot location, with values depicted with short dashed lines. Dotted lines connecting matched knots are shown to guide the eye.

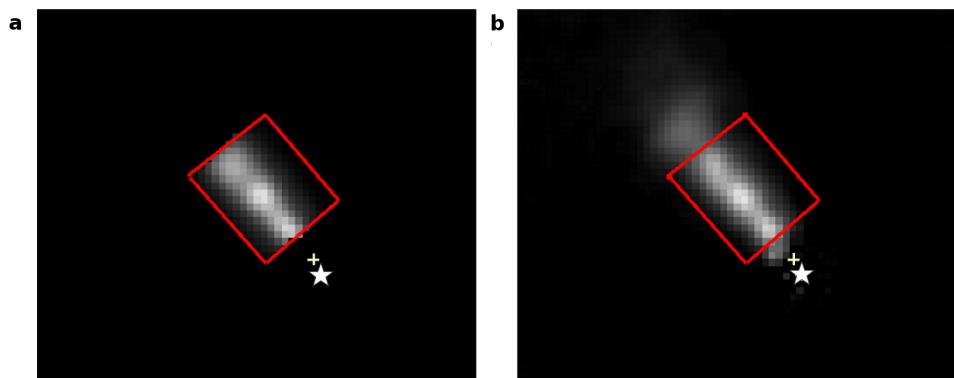


Extended Data Figure 5 | Comparison of 1983 radio positions of knots with HST data. The locations of knots B and C in 1983 are denoted by triangles, and we include an estimated 15 mas error on the position based on a typical centroiding error of 10% of the beam size ($0.15''$). Optical data are identical to that described in Fig. 2. For knot C, a dashed line gives the fit to epochs 1983 and 1994. For knot B, the dashed line is a parabolic fit to all data. The location of knot A is noted with a dotted line.



Extended Data Figure 6 | Depiction of background subtraction method.
a, The final 2014 image stack before subtraction. **b**, The mask used in modelling the galaxy and dust disk. (White areas were not used in fits.) **c**, The final 2014

galaxy and dust model. **d**, The same image as **a** after **c** has been subtracted. Physical scaling of all images and flux scaling of **a**, **b** and **d** are identical.



Extended Data Figure 7 | Depiction of modelling results. **a**, The resulting background-subtracted model image for the 2002 epoch. **b**, The real background-subtracted 2002 image. In both cases the image box is shown by the red outline. Only pixels within this area were used in the fit.

Extended Data Table 1 | Maximum-likelihood modelling results

Name	Year	Radial Position (pixels)	Transverse Position (pixels)	Intensity (μ Jy/pixel)	Radius (pixels)	Radial Position* (mas)	Total Flux (μ Jy)	Size (pc)
Knot A [†]	1994	5.45 ± 0.13	...	0.65 ± 0.04	1.28 ± 0.04	136 ± 5	3.3 ± 0.3	28.3 ± 1.0
	1996	4.95 ± 0.12	...	0.43 ± 0.02	2.01 ± 0.05	124 ± 5	5.5 ± 0.3	44.4 ± 1.0
	2002	4.89 ± 0.09	...	0.24 ± 0.01	2.05 ± 0.04	122 ± 4	3.1 ± 0.1	45.3 ± 0.9
	2014	4.97 ± 0.12	...	0.47 ± 0.02	1.42 ± 0.03	124 ± 3	3.0 ± 0.1	31.4 ± 0.6
Knot B	1994	10.28 ± 0.15	0.45 ± 0.13	0.37 ± 0.02	1.76 ± 0.05	257 ± 5	3.6 ± 0.3	38.8 ± 1.2
	1996	11.26 ± 0.11	-0.32 ± 0.11	0.42 ± 0.02	1.62 ± 0.04	281 ± 5	3.5 ± 0.2	35.9 ± 1.0
	2002	12.11 ± 0.06	0.28 ± 0.05	0.41 ± 0.01	1.79 ± 0.02	303 ± 4	4.2 ± 0.1	39.6 ± 0.5
	2014	14.35 ± 0.07	-0.70 ± 0.05	0.42 ± 0.01	1.98 ± 0.02	359 ± 2	5.3 ± 0.1	43.9 ± 0.5
Knot C	1994	18.02 ± 0.19	-0.08 ± 0.15	0.24 ± 0.02	2.06 ± 0.07	450 ± 6	3.2 ± 0.3	45.6 ± 1.6
	1996	18.39 ± 0.11	-0.41 ± 0.09	0.36 ± 0.01	1.98 ± 0.04	460 ± 4	4.5 ± 0.2	43.7 ± 0.9
	2002	18.32 ± 0.07	-0.28 ± 0.06	0.27 ± 0.01	2.21 ± 0.03	458 ± 4	4.2 ± 0.1	48.9 ± 0.6
	2014	19.14 ± 0.10	-0.56 ± 0.07	0.28 ± 0.01	2.51 ± 0.03	478 ± 2	5.5 ± 0.1	55.4 ± 0.6

* Systematic error of 3.5 (0.18) mas for WFPC2 (ACS) images has been added to the measurement error from column 3. [†]Knot A values are included for completeness; see Methods for note on reliability of knot A positions.

Extreme ultraviolet high-harmonic spectroscopy of solids

T. T. Luu^{1*}, M. Garg^{1*}, S. Yu. Kruchinin¹, A. Moulet¹, M. Th. Hassan¹ & E. Goulielmakis¹

Extreme ultraviolet (EUV) high-harmonic radiation^{1,2} emerging from laser-driven atoms, molecules or plasmas underlies powerful attosecond spectroscopy techniques^{3–5} and provides insight into fundamental structural and dynamic properties of matter^{6,7}. The advancement of these spectroscopy techniques to study strong-field electron dynamics in condensed matter calls for the generation and manipulation of EUV radiation in bulk solids, but this capability has remained beyond the reach of optical sciences. Recent experiments^{8,9} and theoretical predictions^{10–12} paved the way to strong-field physics in solids by demonstrating the generation and optical control of deep ultraviolet radiation⁸ in bulk semiconductors, driven by femtosecond mid-infrared fields or the coherent up-conversion of terahertz fields to multi-octave spectra in the mid-infrared and optical frequencies⁹. Here we demonstrate that thin films of SiO₂ exposed to intense, few-cycle to sub-cycle pulses give rise to wideband coherent EUV radiation extending in energy to about 40 electronvolts. Our study indicates the association of the emitted EUV radiation with intraband currents of multi-petahertz frequency, induced in the lowest conduction band of SiO₂. To demonstrate the applicability of high-harmonic spectroscopy to solids, we exploit the EUV spectra to gain access to fine details of the energy dispersion profile of the conduction band that are as yet inaccessible by photoemission spectroscopy in wide-bandgap dielectrics. In addition, we use the EUV spectra to trace the attosecond control of the intraband electron motion induced by synthesized optical transients. Our work advances lightwave electronics^{5,13–15} in condensed matter into the realm of multi-petahertz frequencies and their attosecond control, and marks the advent of solid-state EUV photonics.

Coherent radiation, emerging from the nonlinear interaction of intense lasers with matter, contains information about structure and dynamics and has lent itself to a wide range of fundamental spectroscopy techniques^{3,16}. In strong-field laser physics³, acceleration of electron wavepackets¹⁷ launched by the strong-field ionization¹⁸ of atoms or molecules and their re-collision with their parent ions gives rise to coherent attosecond EUV radiation^{1,2}. Such radiation can provide detailed information about electronic structure¹⁹, the atomic motion in molecules^{20,21}, and the dynamics of free⁷ and bound²² electron wavepackets in the generated ionic species. Here we demonstrate the extension of strong-field EUV spectroscopy and photonics to the condensed phase.

In our experiments, thin (~ 120 nm), free-standing, polycrystalline SiO₂ films (see Supplementary Information section I) were exposed to intense field-sculpted, ultrafast optical waveforms generated in a second-generation light field synthesizer (refs 23 and 24 and M.Th.H. *et al.*, manuscript in preparation). In Fig. 1c we show spectra of the radiation emerging from the bulk of the SiO₂ samples (Fig. 1b) when driven by the field of a 1.5-cycle pulse (Fig. 1a; energy ~ 0.5 μ J) as a function of the transmitted field amplitude that was varied in the range ~ 0.7 – 1.1 V \AA^{-1} (in the bulk of SiO₂). Regularly spaced spectral peaks, the energy of which match well the odd harmonics of the ultrashort

pulse carrier frequency, attest to the coherence between the driving and emitted EUV fields.

A gradual increase in the field strength of the optical pulse gives rise to the formation of a robust spectral plateau extending over more than ~ 20 eV (15–35 eV) in the EUV range. The cutoff energy, defined as the last harmonic peak discernible at the high-energy end of the EUV spectral plateau, undergoes a linear increase with optical field strength (Fig. 1d). The intensity yield of the emitted harmonic radiation exhibits a highly nonlinear dependence on optical bias, unique for each individual harmonic peak (Fig. 1e–g). The slope of the intensity yield declines at higher optical biasing of the medium (Fig. 1e), while for some of the harmonic peaks, it rises again with further increase of the field strength (Fig. 1f, g).

The linear scaling of the harmonic cutoff energy with the field amplitude distinguishes the underlying mechanism of high-harmonic radiation from that in a gas phase¹⁷, agrees with the scaling observed in previous studies of coherent, laser-driven dynamics in solids giving rise to radiation in the terahertz²⁵, optical⁹ and vacuum ultraviolet⁸ ranges, and dictates the need for theoretical modelling tailored to describe the interaction of intense optical fields with bulk solids. To gain insight into the physical dynamics underlying EUV emission in our experiments, we conducted both semiclassical¹⁰ and quantum-mechanical simulations¹¹ under the precise experimental conditions given above.

In the tight-binding model beyond the nearest-neighbour approximation, the dependence of energy $E_i(k)$ in the i th band on crystal momentum k is represented by a series of spatial harmonics:

$$E_i(k) = \sum_{n=0}^{n_{\max}} \varepsilon_{i,n} \cos(nka) \quad (1)$$

Here a is the lattice constant, n is the lattice site number (order of the spatial harmonic), n_{\max} is the maximum number of distant neighbours included in the model, $\varepsilon_{i,n}$ are Fourier coefficients, and $\varepsilon_{i,0}$ is the band offset (see Supplementary Information section II). The highest characteristic distance in the crystal is, therefore, $R_{\max} = n_{\max}a$. Consequently, the maximum energy that an electron wavepacket can acquire from the field is $E_{\max} = eF_0R_{\max} = n_{\max}\hbar\omega_B$, where $\omega_B = eF_0a/\hbar$ is the Bloch frequency and F_0 is the field amplitude.

To illustrate the role of the energy dispersion of the lowest conduction band in the emission of EUV radiation, we first perform semiclassical simulations of a pre-existing wavepacket driven by the electric field in Fig. 1a, for $F_0 = 0.8$ V \AA^{-1} along the contour of different spatial components, which are included in the Fourier series (equation (1)) of the band dispersion (Fig. 2b); see Supplementary Information section VIII for the case of $F_0 = 1.1$ V \AA^{-1} . In our simulations, the amplitudes $\varepsilon_{c,n}$ as well as the number of significant spatial harmonics, are derived from the Fourier expansion of the band along the Γ –M direction in the Brillouin zone (see Supplementary Information section II). The simulations (Fig. 1d) highlight the gradual increase of the cutoff energy with spatial harmonic order n and suggest that, under our experimental conditions, deceleration of electrons by the spatial harmonics of orders

¹Max-Planck-Institut für Quantenoptik, Hans-Kopfermann-Strasse 1, D-85748 Garching, Germany.

*These authors contributed equally to this work.

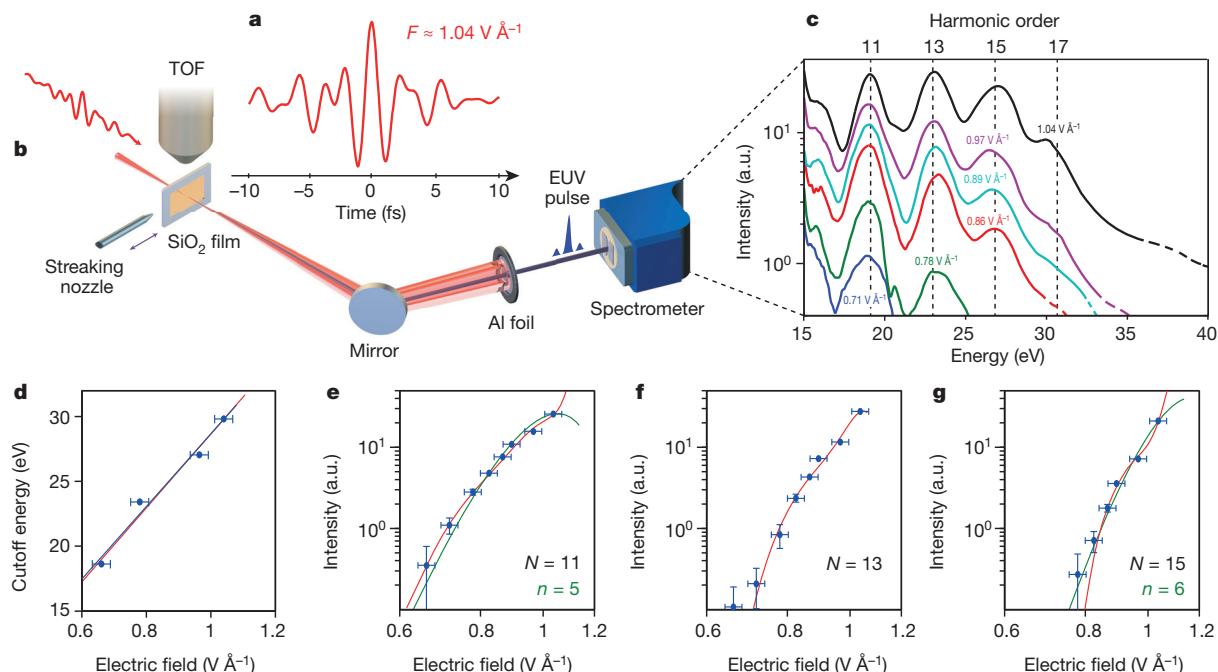


Figure 1 | Generation, energy cutoff and intensity scaling of coherent EUV radiation in SiO_2 films. a, The 1.5-cycle driving pulse (with spectral range $\Delta_L \approx 1.1\text{--}2.5\text{ eV}$) produced in a light field synthesizer and characterized by attosecond streaking, before being focused onto a thin film of SiO_2 ($\sim 120\text{ nm}$) to generate EUV radiation (b). c, EUV spectra recorded for increasing field strengths for which the screening factor (~ 0.8) derived from the Fresnel formula is taken into account. The dashed portion of the curves indicates the noise floor. d, Scaling of the cutoff energy E_{cutoff} (blue dots with error bars) with the field amplitude F_0 . Shown is linear fitting of the measured data (blue line) and results of semiclassical simulations

(red line). e–g, Yield dependence for harmonics of order $N = 11$ (e), 13 (f) and 15 (g) on the field amplitude (blue dots). Red curves represent fittings based on equation (3), including all spatial harmonics up to n_{max} . Green curves in e, g, represent the fitting of experimental data with equation (3), where only one spatial harmonic ($n = 5$ and $n = 6$, respectively) was taken into account. Error bars depict the standard deviations of the values extracted from several data sets recorded under identical conditions. Corresponding slopes of the intensity curves at $0.7\text{ V } \text{\AA}^{-1}$ are evaluated as $\sim F_0^{10}$, $\sim F_0^{18}$, and $\sim F_0^{15}$, respectively.

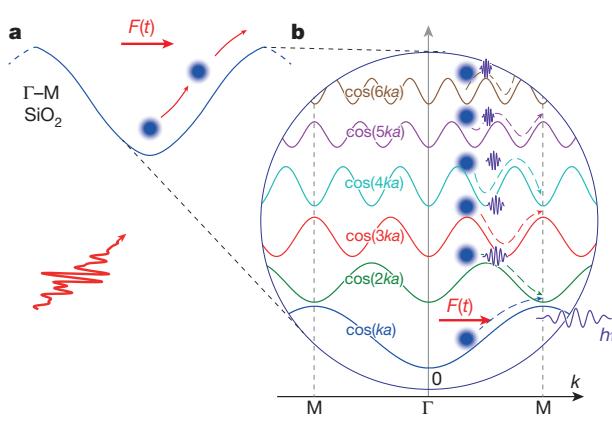
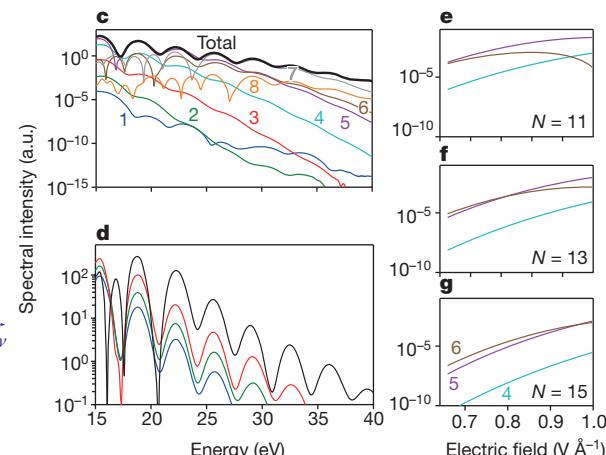


Figure 2 | Semiclassical picture of the field-driven electron dynamics in SiO_2 . a, An electron wavepacket driven by the field in the lowest conduction band along the $\Gamma\text{--}M$ direction of the Brillouin zone. b, Spatial harmonic components of the band dispersion resulting from equation (15) in Supplementary Information. Relative amplitudes are scaled, for clarity. c, Contributions to the EUV spectra from each of the spatial harmonics for the field of Fig. 1a at the amplitude $F_0 = 0.8\text{ V } \text{\AA}^{-1}$. The total emitted spectrum



(black curve) aggregates contributions from all spatial harmonics. d, Simulated spectra for the electric field waveform (Fig. 1a) for the realistic band shown in a as a function of the peak field strength F_0 . Colour-coding corresponds to field strengths of Fig. 1c. e–g, Contributions of spatial harmonics $n = 4, 5, 6$ to the intensity yields of EUV harmonics of orders $N = 11$ (e), 13 (f) and 15 (g) versus field amplitude F_0 .

$n = 5, 6$, and 7 play a central role in the emission of radiation near the cutoff energy.

From the linear fitting of experimental data (Fig. 1d, blue line), we evaluate the largest characteristic distance significantly contributing to the spectral intensity as $R_{\text{cutoff}} = n_{\text{cutoff}}a = e^{-1}dE_{\text{cutoff}}/dF_0$, and obtain

$R_{\text{cutoff}} \approx 30\text{ \AA}$ and $n_{\text{cutoff}} = 6.2 \pm 0.5$. R_{cutoff} is smaller than the average size of the crystalline domain (which is $> 40\text{ \AA}$), and thus, a scattering at domain edges can be neglected (see Supplementary Information section I). Extrapolation of the linear fitting to the zero field $F_0 = 0$ yields a cutoff energy $E_{\text{cutoff}} = 0.6 \pm 2.0\text{ eV}$ (see Supplementary

Information section III) and even further strengthens the link between the emitted radiation and the induction of intraband currents considered in the semiclassical treatment, where $E_{\text{cutoff}} = 0$ for $F_0 = 0$ is predicted. When performed under the realistic energy dispersion of the lowest conduction band $E_c(k)$ along the Γ -M direction (Fig. 2a), the semiclassical simulation yields a reasonable reproduction (Fig. 2d) of the experimental spectra shown in Fig. 1c. The calculated maximum cutoff energy $E_{\text{cutoff}} \approx 33$ eV (within the same dynamic range of our experiments), its linear scaling versus field amplitude F_0 , and $n_{\text{cutoff}} = 6$ evaluated from the slope of the cutoff dependence of the simulated spectra (red line, Fig. 1d) are in excellent agreement with our experimental findings (blue line, Fig. 1d). Simulations assuming that the field polarization is parallel to the Γ -K or Γ -A directions of SiO_2 yielded EUV spectra of an order of magnitude weaker intensity and lower cutoff energy in comparison to those along the Γ -M direction (see Supplementary Information section II, Supplementary Fig. 2). The contribution of the highest valence bands was also found to be negligible (see Supplementary Information section II). The above conclusions support the notion that, in our experiments, the EUV emission primarily occurs due to acceleration of first conduction band electrons in microcrystals, for which the Γ -M direction is parallel to the laser polarization axis.

In our quantum-mechanical simulations based on the numerical solution of the semiconductor Bloch equations^{9,11}, we find that the cutoff energy and fine details of our experimental spectra are reasonably reproduced by the intraband current only, but they disagree with those obtained from interband or total polarization (see Supplementary Information section IV, Supplementary Figs 8, 9). These results were further verified by contrasting the intensity dependence of the 11th harmonic in our experiment with that predicted by the semiclassical model and interband and intraband terms of quantum-mechanical simulations (see Supplementary Information section V).

The previous comparison of simulations with measurements pinpoints the importance of intraband dynamics in describing the process and the capability of the semiclassical approach to capture key aspects of the emission of EUV radiation, including the nonlinearity pertaining to the intensity scaling. As a result, the characteristics of the EUV spectra offer direct insight into the generation of multi-PHz field-tunable intraband currents, whose cutoff frequency is adjustable from $v_{\text{cutoff}} \approx 3.6$ PHz to 8.5 PHz (Fig. 1c). These are the highest frequency currents so far induced in a solid medium by optical fields, and more than an order of magnitude higher than those in previous experiments in solid-state lightwave electronics^{13–15}.

The semiclassical theory of high-harmonic radiation²⁶ predicts that the intensity of the emitted spectra in the case of a cosine band exhibits an extreme nonlinearity versus the field amplitude F_0 . For an electric field carried at a frequency ω_L , the intensity dependence of the Nth order harmonic is proportional to

$$I_N \propto (N\omega_L)^2 \left| J_N \left[\frac{eF_0a}{\hbar\omega_L} \right] \right|^2 \quad (2)$$

where $J_N(z)$ is the Bessel function of the first kind of order N .

For a band whose dispersion profile comprises a few spatial harmonics, equation (2) can be generalized (see Supplementary Information section II) to account for the contributions of distant neighbours to the yield of the Nth harmonic peak as

$$I_{c,N} \propto (N\omega_L)^2 \left| \sum_{n=1}^{n_{\text{max}}} n e_{c,n} J_N \left[\frac{n\omega_B}{\omega_L} \right] \right|^2 \quad (3)$$

These equations are also accurate in the short pulse limit with the difference that, in this case, F_0 refers to the peak field strength of the pulse. Figure 2e–g illustrates how the representative spatial harmonics ($n = 4, 5, 6$) influence the yield dependence of EUV harmonic peaks in our spectral range versus F_0 according to equation (3). These plots underpin the possibility of extracting structural information about the

underlying medium by fitting the yield of a single or several EUV harmonic peaks by equation (3). Indeed, a simultaneous fitting for harmonic orders 11 to 15 (see Fig. 1e–g, red curves) by equation (3) yields an excellent regression and allows for retrieval of the amplitudes $e_{c,n}$ in expansion (1).

Based on the confidence interval of the reconstruction, our fitting reveals that the intensity dependence of EUV harmonics in the spectral range of our study aggregates primarily the contributions of EUV photons emitted when the field-driven electron wavepacket is decelerated by the spatial harmonics $n = 5$ and 6. For instance, the yield of EUV harmonic $N = 11$ is dominated by the deceleration of the wavepacket by the spatial harmonic $n = 5$ (Fig. 1e, green curve), while for $N = 15$, deceleration by $n = 6$ dominates (Fig. 1g, green curve); we verify this by fitting their yield based on equation (3) for the corresponding spatial harmonics. This result is in excellent agreement with the independent conclusions of the cutoff analysis of Fig. 1d, which reveals similar sensitivity to spatial harmonics of the cutoff energy of the emitted radiation.

The evaluated ratio of the significantly contributing coefficients is $e_{c,5}/e_{c,6} \approx -5.1 \pm 0.6$ and accurately matches the theoretically predicted ratio²⁷ obtained in the expansion given by equation (15) in the Supplementary Information, $e_{c,5}/e_{c,6} \approx -5.6$. This result offers additional evidence that the laser-driven intraband current accurately captures the EUV emission and, most importantly, constitutes the first measurement of the fine details of the conduction band dispersion in wide bandgap materials. Indeed, the large insulating bandgap and low mobility of photo-excited carriers prevent replenishing of the created vacancies, give rise to charging effects, which in turn make photoemission spectroscopies²⁸ impractical, and limit their reach to valence or low-lying conduction bands²⁸. As a pure polarization spectroscopy, our approach overcomes this essential frontier and highlights the promise for a new range of studies in dielectric materials. The resolution pertinent to our technique can be estimated from the highest spatial harmonic order influencing the EUV emission as $\pi/6a \approx 0.1 \text{ \AA}^{-1}$ (for $n_{\text{cutoff}} = 6$), which is close to that of modern angular resolved photoemission spectroscopy.

In a second set of experiments, we dynamically probe the physics underlying EUV generation in the condensed phase and explore routes towards the ultrafast control of multi-PHz-scale currents by conducting measurements of EUV spectra generated by various waveforms including few-cycle, 1.5-cycle, sub-cycle and half-cycle pulses. Spectra recorded for two settings of the carrier envelope phase (CEP) of the above driving fields are shown in Fig. 3. The increase of the averaged inter-peak energy spacing (Fig. 3a to Fig. 3b) and the broadening of these harmonic peaks (Fig. 3c) are compatible with increasing driver frequencies and decreasing pulse duration of the optical field, respectively, and further corroborate the coherent relation between the emitted and driving fields. In contrast to few-cycle and 1.5-cycle driving fields (Fig. 3a, b), sub-cycle and, even more evidently, half-cycle field waveforms give rise to EUV continua (Fig. 3c, d), whose amplitude and shape show a considerable sensitivity to the CEP adjustments.

The dynamic control of the emission attained by CEP manipulation of the waveform of the driver pulses lends itself to use as an additional platform to verify the validity of the conclusions reached in the cutoff and intensity analysis (Fig. 1d–g), by further contrasting our findings with quantum-mechanical and semiclassical models. Figure 3e, f, k, l shows simulated temporal dynamics of an electron wavepacket in the conduction band of SiO_2 for two representative waveforms of the few-cycle and half-cycle driving fields, and for the phase settings of the insets of Fig. 3a, d. Interband dynamics give rise to periodic modulations of the amplitude of the wavepacket as a function of time, whereas the intraband dynamics are best manifested by the lateral oscillatory displacement of the excited wavepacket. Figure 3g, m depicts the simulated temporal intensity profiles of the radiation induced by intraband currents, spectrally filtered in the energy range of our experiments (>15 eV); Fig. 3h, n shows the results of semiclassical simulations;

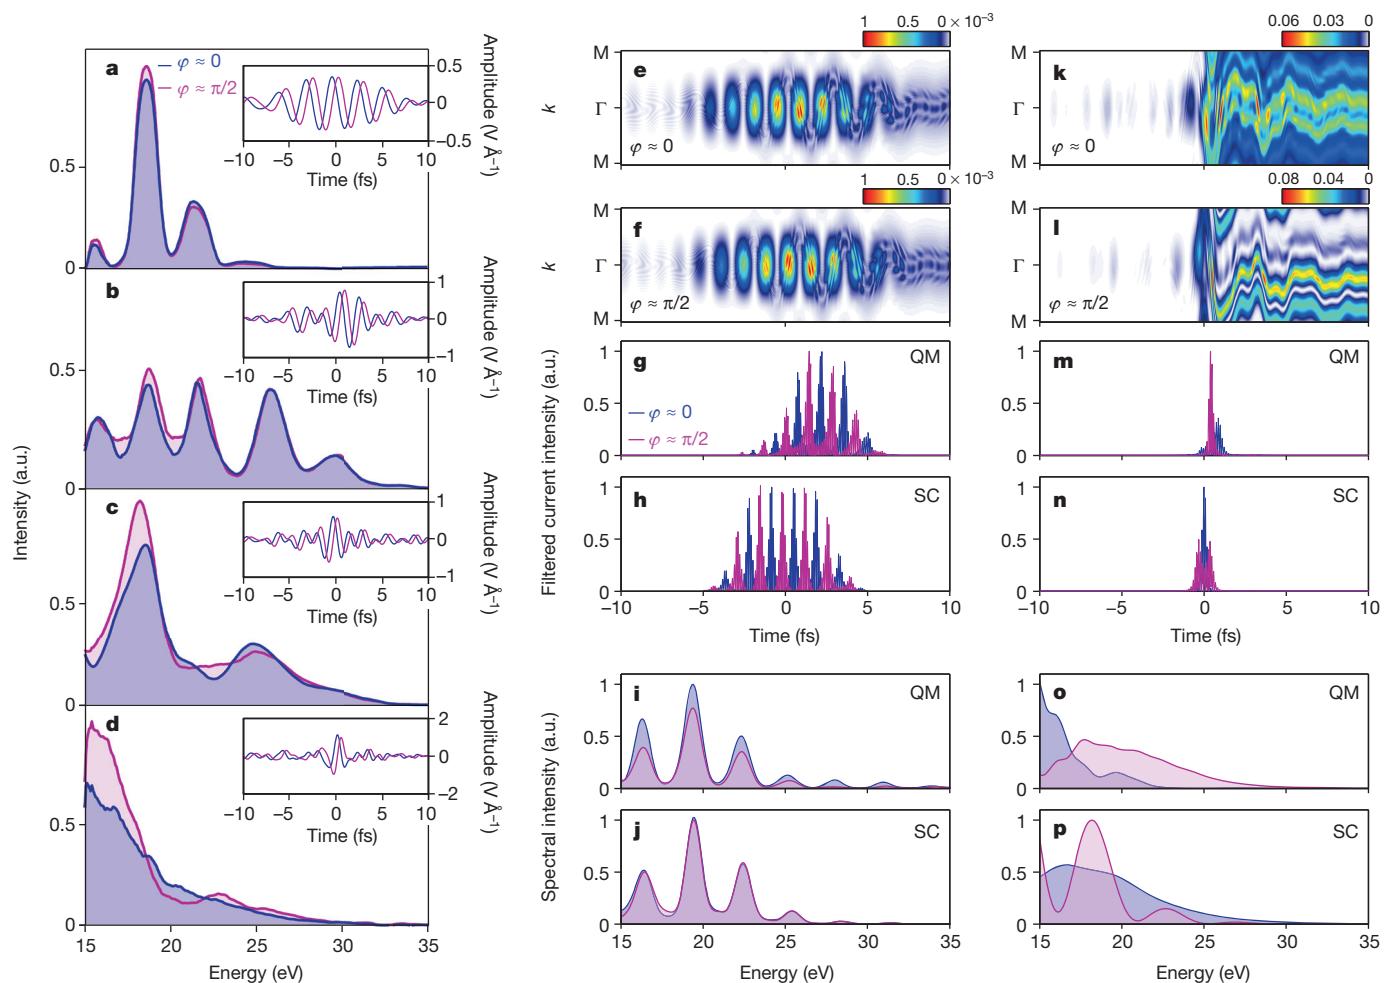


Figure 3 | Field control of laser-driven EUV emission in SiO_2 (experiment and simulation). **a–d**, EUV harmonic spectra recorded for two CEP settings ($\varphi_{CE} \approx 0$, blue; $\varphi_{CE} \approx \pi/2$ rad, magenta) of the driving waveforms shown in insets: **a**, few-cycle ($\mathcal{A}_L \approx 1.1\text{--}1.8$ eV and carrier frequency $\omega_L \approx 1.52$ eV), **b**, 1.5-cycle ($\mathcal{A}_L \approx 1.1\text{--}2.5$ eV, $\omega_L \approx 1.96$ eV), **c**, sub-cycle ($\mathcal{A}_L \approx 1.1\text{--}3.1$ eV, $\omega_L \approx 2.227$ eV), and **d**, half-cycle ($\mathcal{A}_L \approx 1.1\text{--}4.6$ eV, $\omega_L \approx 2.1$ eV) pulses. The waveforms are sampled by attosecond streaking. **e–f**, Population dynamics in the lowest conduction band of SiO_2 along the Γ -M direction driven by the few-

cycle pulse with $\varphi_{CE} \approx 0$ and $\varphi_{CE} \approx \pi/2$ rad, respectively (insets of **a**, **g**, **h**). Temporal intensity profiles $I(t) \propto |\partial J(t)/\partial t|^2$ of the radiated EUV fields in the spectral region > 15 eV for a few-cycle pulse, calculated with quantum-mechanical (QM) and semiclassical (SC) models. $J(t)$ denotes the intraband current. **i**, Corresponding spectra obtained from quantum-mechanical and **j**, semiclassical simulations for a few-cycle pulse for the above-mentioned CEP settings are marked by the same colour code. **k–p**, Same as **e–j** but calculated for a half-cycle pulse (insets of **d**).

and Fig. 3*i, j, o, p* show the corresponding spectra (see Supplementary Information section IV for additional information). Importantly and in support of the conclusions of the intensity scaling analysis, the results of the semiclassical simulations for the entire range of pulse waveforms exploited are virtually identical, if we consider the band dispersion with the terms $\sim \cos(5ka) - 0.19\cos(6ka)$, which represent significantly contributing spatial harmonics and their experimentally retrieved relative amplitudes (Fig. 1*d–g*).

The reasonable reproduction of the experimental spectra for the broad range of waveforms used (see also Supplementary Information section IV) by both the semiclassically and quantum-mechanically simulated intraband currents allows the linking of the temporal electron dynamics underlying our simulations with those in our experiments. For few-cycle pulses (Fig. 3*g*), excitation of the currents and concomitant radiation extend temporally over several field cycles. The current profile (Fig. 3*g, h*, red and blue curves) in this case is temporally displaced by a CEP variation of the optical driving field, but its structure remains virtually unaffected, resulting in nearly invariable spectra (Fig. 3*i, j*) with respect to the variation of this pulse parameter, in agreement with our experiments (Fig. 3*a*). By contrast, half-cycle pulses confine and control electric currents to the optical cycle

(Fig. 3*m, n*) with CEP. This control is manifested in the spectral domain by wideband modulations of the broadband emitted spectra (Fig. 3*o, p*) and their continuum shape to give structured ones in accordance with our experiments (Fig. 3*d*). Notably, the semiclassical simulation (Fig. 3*j, p*) seems to reproduce this trend better than the quantum-mechanical one (Fig. 3*i, o*). This can be attributed to the fact that intraband current and interband polarization are inherently coupled in the semiconductor Bloch equations.

The CEP-based manipulation of the emitted spectra emanating from an extreme nonlinear process has served for years as a paradigm of sub-cycle confinement and control in attosecond physics^{6,13}. Our experiments and simulations support the notion that this paradigm is also applicable for strong-field-driven electron dynamics and EUV emission in bulk solids. As half-cycle drivers are confined to a *de facto* sub-femtosecond interval (duration $\tau_p < 1$ fs), these broadband spectral modulations offer evidence of the attosecond control of the induced multi-PHz intraband currents, and represent a unique and generic tool for tracing coherent strong-field dynamics in solids.

Further study of the coherence of the emitted EUV radiation included the measurement of the far-field spatial properties of the emerging radiation. The measurements reveal a well-behaved (ellipti-

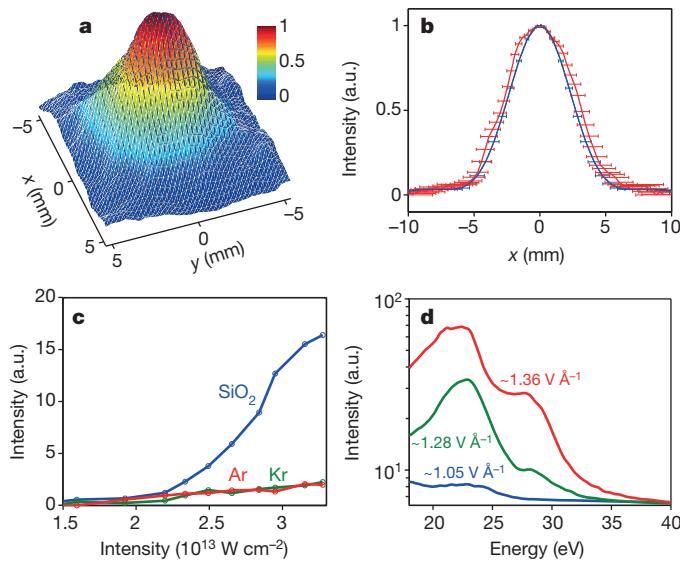


Figure 4 | Coherent EUV source characteristics. **a**, Far-field spatial intensity profile of the EUV beam sampled by a detector placed at a distance of ~ 90 cm downstream from the source. **b**, x -axis profile (red curve) of **a**, and far-field EUV beam profile (blue curve) calculated with semiclassical simulations and free-space propagation. Error bars depict the standard deviations of the values extracted from several data sets recorded under identical conditions. **c**, EUV photon yield (1.5-cycle pulse) integrated over the range (15–40 eV) for SiO_2 (blue curve) and for the noble gases Ar (red curve) and Kr (green curve). **d**, EUV supercontinua in SiO_2 extending to ~ 40 eV driven by half-cycle pulses at different peak field strengths (given on curves).

city ~ 0.93) EUV beam profile (Fig. 4a), whose dimensions and shape are in excellent agreement with those predicted by our semiclassical simulations and free-space propagation (Fig. 4b, blue curve), and suggest the full spatial coherence of the emitted EUV (see Supplementary Information section VI).

To gain insights into the conversion efficiency of laser radiation to extreme ultraviolet, we compared the yield of the EUV emission in SiO_2 with that of two broadly used media in gas-phase attosecond physics, Kr and Ar (Fig. 4c), under identical conditions (see Supplementary Information section VII). Our study reveals that the yield of EUV emission from SiO_2 , which is approximately 8 times higher than that of these gases in the range of applied optical fields (Fig. 1c), falls within a similar efficiency class (10^{-7} – 10^{-6}) as that of gases²⁹.

Extension of the EUV continua to higher photon energies (up to 40 eV or $\nu \approx 9.7$ PHz) is also possible by applying stronger field ($\sim 1.36 \text{ V \AA}^{-1}$) to our samples (Fig. 4d). For an even stronger optical field, the ultrathin samples show degradation due to physical melting. However, tests on thicker samples showed suppression of melting effects and may soon enable advancement of the coherent emission towards the soft X-ray regime, catalysing further developments.

The generation of broadband EUV radiation by driving the coherent nonlinear motion of electrons in bulk solids advances attosecond high-harmonic spectroscopy to the condensed phase, and promotes solid-state laser photonics and electronics to a new frequency regime with great potential for novel scientific and technological inquiry. It offers new possibilities for the tracing and attosecond control of strong-field electron dynamics in solids.

Received 22 June 2014; accepted 27 March 2015.

- Ferry, M. et al. Multiple-harmonic conversion of 1064 nm radiation in rare gases. *J. Phys. B* **21**, L31–L35 (1988).
- L'Huillier, A. & Balcou, P. High-order harmonic generation in rare gases with a 1-ps 1053-nm laser. *Phys. Rev. Lett.* **70**, 774–777 (1993).
- Corkum, P. B. & Krausz, F. Attosecond science. *Nature Phys.* **3**, 381–387 (2007).
- Hentschel, M. et al. Attosecond metrology. *Nature* **414**, 509–513 (2001).
- Goulielmakis, E. et al. Attosecond control and measurement: lightwave electronics. *Science* **317**, 769–775 (2007).
- Baltuška, A. et al. Attosecond control of electronic processes by intense light fields. *Nature* **421**, 611–615 (2003).
- Remetter, T. et al. Attosecond electron wave packet interferometry. *Nature Phys.* **2**, 323–326 (2006).
- Ghimire, S. et al. Observation of high-order harmonic generation in a bulk crystal. *Nature Phys.* **7**, 138–141 (2011).
- Schubert, O. et al. Sub-cycle control of terahertz high-harmonic generation by dynamical Bloch oscillations. *Nature Photon.* **8**, 119–123 (2014).
- Feise, M. W. & Citrin, D. S. Semiclassical theory of terahertz multiple-harmonic generation in semiconductor superlattices. *Appl. Phys. Lett.* **75**, 3536–3538 (1999).
- Golde, D., Meier, T. & Koch, S. W. High harmonics generated in semiconductor nanostructures by the coupled dynamics of optical inter- and intraband excitations. *Phys. Rev. B* **77**, 075330 (2008).
- Mücke, O. D. Isolated high-order harmonics pulse from two-color-driven Bloch oscillations in bulk semiconductors. *Phys. Rev. B* **84**, 081202 (2011).
- Krüger, M., Schenk, M. & Hommelhoff, P. Attosecond control of electrons emitted from a nanoscale metal tip. *Nature* **475**, 78–81 (2011).
- Schiffelin, A. et al. Optical-field-induced current in dielectrics. *Nature* **493**, 70–74 (2013).
- Schultze, M. et al. Controlling dielectrics with the electric field of light. *Nature* **493**, 75–78 (2013).
- Levenson, M. D. & Kano, S. *Introduction to Nonlinear Laser Spectroscopy* Revised edn (Academic, 1988).
- Corkum, P. B. Plasma perspective on strong-field multiphoton ionization. *Phys. Rev. Lett.* **71**, 1994–1997 (1993).
- Keldysh, L. V. Ionization in field of a strong electromagnetic wave. *Sov. Phys. JETP* **20**, 1307–1314 (1965).
- Itatani, J. et al. Tomographic imaging of molecular orbitals. *Nature* **432**, 867–871 (2004).
- Baker, S. et al. Probing proton dynamics in molecules on an attosecond time scale. *Science* **312**, 424–427 (2006).
- Wörner, H. J., Bertrand, J. B., Kartashov, D. V., Corkum, P. B. & Villeneuve, D. M. Following a chemical reaction using high-harmonic interferometry. *Nature* **466**, 604–607 (2010).
- Smirnova, O. et al. High harmonic interferometry of multi-electron dynamics in molecules. *Nature* **460**, 972–977 (2009).
- Hassan, M. T. et al. Attosecond photonics: synthesis and control of light transients. *Rev. Sci. Instrum.* **83**, 111301 (2012).
- Wirth, A. et al. Synthesized light transients. *Science* **334**, 195–200 (2011).
- Waschke, C. et al. Coherent submillimeter-wave emission from Bloch oscillations in a semiconductor superlattice. *Phys. Rev. Lett.* **70**, 3319–3322 (1993).
- Wegener, M. *Extreme Nonlinear Optics: An Introduction* (Springer, 2005).
- Chelikowsky, J. R. & Schlüter, M. Electron states in α -quartz: a self-consistent pseudopotential calculation. *Phys. Rev. B* **15**, 4020–4029 (1977).
- Damascelli, A. Probing the electronic structure of complex systems by ARPES. *Phys. Scr. T* **109**, 61–74 (2004).
- Cabasse, A., Machinet, G., Dubrouil, A., Cormier, E. & Constant, E. Optimization and phase matching of fiber-laser-driven high-order harmonic generation at high repetition rate. *Opt. Lett.* **37**, 4618–4620 (2012).

Supplementary Information is available in the online version of the paper.

Acknowledgements We thank M. Wismer for support in calculations, and A. Jain for help during experiments. This work was supported by a European Research Council grant (Attoelectronics-258501), the Deutsche Forschungsgemeinschaft Cluster of Excellence: Munich Centre for Advanced Photonics (www.munich-photonics.de), the Max Planck Society and the European Research Training Network ATTOFEL.

Author Contributions T.T.L. and M.G. conducted the experiments; A.M. and M.Th.H. contributed to the development of the source; E.G. conceived the experiments; T.T.L., M.G. and S.Yu.K. conducted the simulations; S.Yu.K. performed the analytical derivations; E.G., T.T.L., M.G. and S.Yu.K. contributed to the preparation of the manuscript.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to E.G. (elgo@mpq.mpg.de).

Robots that can adapt like animals

Antoine Cully^{1,2}, Jeff Clune³, Danesh Tarapore^{1,2,†} & Jean-Baptiste Mouret^{1,2,4,5,6,‡}

Robots have transformed many industries, most notably manufacturing¹, and have the power to deliver tremendous benefits to society, such as in search and rescue², disaster response³, health care⁴ and transportation⁵. They are also invaluable tools for scientific exploration in environments inaccessible to humans, from distant planets⁶ to deep oceans⁷. A major obstacle to their widespread adoption in more complex environments outside factories is their fragility^{6,8}. Whereas animals can quickly adapt to injuries, current robots cannot ‘think outside the box’ to find a compensatory behaviour when they are damaged: they are limited to their pre-specified self-sensing abilities, can diagnose only anticipated failure modes⁹, and require a pre-programmed contingency plan for every type of potential damage, an impracticality for complex robots^{6,8}. A promising approach to reducing robot fragility involves having robots learn appropriate behaviours in response to damage^{10,11}, but current techniques are slow even with small, constrained search spaces¹². Here we introduce an intelligent trial-and-error algorithm that allows robots to adapt to damage in less than two minutes in large search spaces without requiring self-diagnosis or pre-specified contingency plans. Before the robot is deployed, it uses a novel technique to create a detailed map of the space of high-performing behaviours. This map represents the robot’s prior knowledge about what behaviours it can perform and their value. When the robot is damaged, it uses this prior knowledge to guide a trial-and-error learning algorithm that conducts intelligent experiments to rapidly discover a behaviour that compensates for the damage. Experiments reveal successful adaptations for a legged robot injured in five different ways, including damaged, broken, and missing legs, and for a robotic arm with joints broken in 14 different ways. This new algorithm will enable more robust, effective, autonomous robots, and may shed light on the principles that animals use to adapt to injury.

Current damage recovery in deployed robots typically involves two phases: self-diagnosis, followed by selection of the best pre-designed contingency plan^{9,13–15}. Such self-diagnosing robots are expensive, because self-monitoring sensors are expensive, and are difficult to design, because robot engineers cannot foresee every possible situation: this approach often fails either because the diagnosis is incorrect^{13,14} or because an appropriate contingency plan is not provided¹⁵.

Injured animals respond differently (see Fig. 1a): they learn by trial and error how to compensate for damage (for example, learning which limp minimizes pain)^{16,17}. Similarly, trial-and-error learning algorithms could allow robots to creatively discover compensatory behaviours without being limited to their designers’ assumptions about how damage may occur and how to compensate for each type of damage. However, state-of-the-art learning algorithms are impractical because of the “curse of dimensionality”¹²: the fastest algorithms, from fields like reinforcement learning¹² and modular robotics^{10,11}, depend on human demonstrations¹² or take 15 min or more even for relatively small search spaces (for example, 6–12 parameters, requiring 15–30 min^{10,11}). Algorithms without these limitations take several hours¹². Damage recovery would be much more practical and effective if robots

adapted as creatively and quickly as animals do (for example, in less than 2 min) in larger search spaces and without expensive, self-diagnosing sensors.

Here we show that rapid adaptation can be achieved by guiding an intelligent trial-and-error learning algorithm with an automatically generated, pre-computed behaviour–performance map that predicts the performance of thousands of different behaviours (Supplementary Video 1). Current learning algorithms either start with no knowledge of the search space¹² or with minimal knowledge from a few human demonstrations^{12,18}. Our hypothesis is that animals understand the space of possible behaviours and their value from previous experience¹⁹, and that animals adapt by intelligently selecting tests that validate or invalidate whole families of promising compensatory behaviours. The key insight here is that robots could do the same.

Our robots store knowledge from previous experience in the form of a map of the behaviour–performance space. Guided by this map, a damaged robot tries different types of behaviours that are predicted to perform well and, as tests are conducted, updates its estimates of the performance of those types of behaviours. The process ends when the robot predicts that the most effective behaviour has already been discovered. A key assumption is that information about many different behaviours of the undamaged robot will still be useful after damage, because some of these behaviours will still be functional despite the damage. The results of our experiments support this assumption for all the types of damage we tested, revealing that a robot can quickly discover a way to compensate for damage (for example, see Fig. 1c) without a detailed mechanistic understanding of its cause, as occurs with animals. We call this approach ‘Intelligent Trial and Error’ (IT&E) (see Fig. 1d).

The behaviour–performance map is created using a novel algorithm and a simulation of the robot, which can be either a standard physics simulator or automatically discovered¹⁴. The robot’s designers need only describe the dimensions of the space of possible behaviours and a performance measure. For instance, walking gaits could be described by how much each leg touches the ground (a behavioural measure) and speed (a performance measure). An alternative gait behavioural measure could be the percentage of time a robot’s torso has positive pitch, roll, and yaw angles. For grasping, performance could be measured by the amount of surface contact, and it has been demonstrated that 90% of effective poses for the 21-degree-of-freedom human hand can be captured by a three-dimensional behavioural space describing the principal components of ways in which hand-poses commonly vary²⁰. To fill in the behaviour–performance map, an optimization algorithm simultaneously searches for a high-performing solution at each point in the behavioural space (Fig. 2a, b and Extended Data Fig. 1). This step requires simulating millions of behaviours, but needs to be performed only once per robot design before deployment (Supplementary Methods).

A low confidence is assigned to the predicted performance of behaviours stored in this behaviour–performance map because they have not been tried in reality (Fig. 2b and Extended Data Fig. 1). During the robot’s mission, if performance drops below a user-defined threshold (owing either to damage or to a different environment), the robot

¹Sorbonne Universités, Université Pierre et Marie Curie (UPMC), Paris 06, UMR 7222, Institut des Systèmes Intelligents et de Robotique (ISIR), F-75005, Paris, France. ²CNRS, UMR 7222, Institut des Systèmes Intelligents et de Robotique (ISIR), F-75005, Paris, France. ³Department of Computer Science, University of Wyoming, Laramie, Wyoming 82071, USA. ⁴Inria, Team Larsen, Villers-lès-Nancy, F-54600, France. ⁵CNRS, Loria, UMR 7503, Vandœuvre-lès-Nancy, F-54500, France. ⁶Université de Lorraine, Loria, UMR 7503, Vandœuvre-lès-Nancy, F-54500, France. [†]Present addresses: Department of Electronics, University of York, York YO10 5DD, UK (D.T.); Inria, Villers-lès-Nancy, F-54600, France (J.-B.M.).

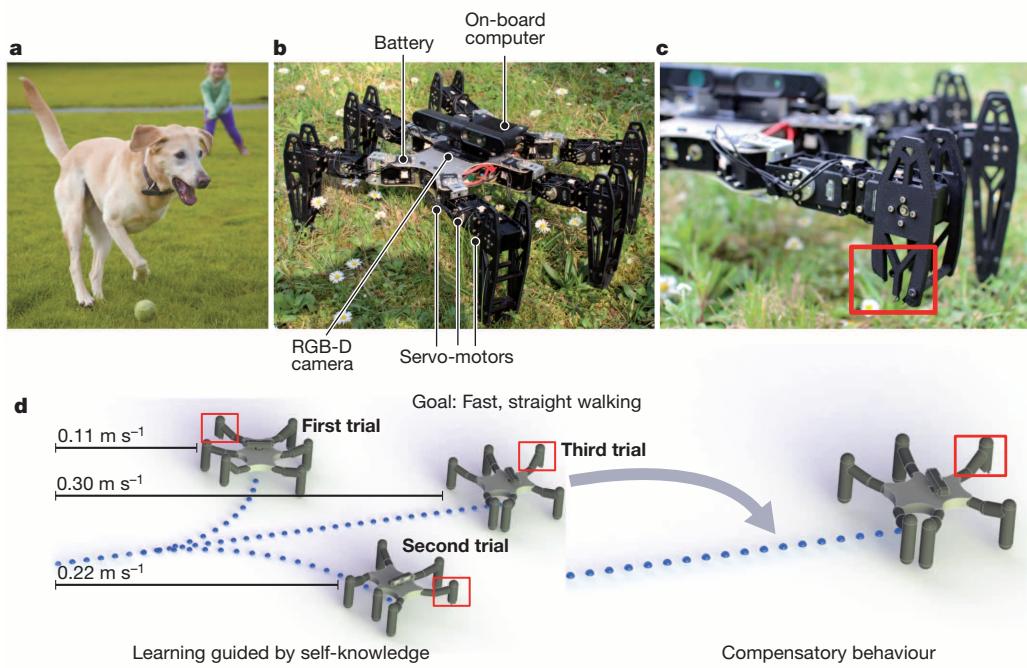


Figure 1 | Using the Intelligent Trial and Error (IT&E) algorithm, robots, like animals, can quickly adapt to recover from damage. **a**, Most animals can find a compensatory behaviour after an injury. Without relying on predefined compensatory behaviours, they learn how to avoid behaviours that are painful or no longer effective. Photo credit: Michael Lloyd/*The Oregonian* (2012). **b**, An undamaged, hexapod robot. RGB-D stands for red, green, blue, depth. **c**, One type of damage the hexapod may have to cope with

selects the most promising behaviour from the behaviour–performance map, tests it, and measures its performance. The robot subsequently updates its prediction for that behaviour and nearby behaviours, assigns high confidence to these predictions (Fig. 2c, d and Extended Data Fig. 1), and continues the select–test–update process until it finds a satisfactory compensatory behaviour (Fig. 2e and Extended Data Fig. 1).

All of these ideas are technically captured via a Gaussian process model²¹, which approximates the performance function with already acquired data, and a Bayesian optimization procedure^{22,23}, which exploits this model to search for the maximum of the performance function (see Supplementary Methods). The robot selects which behaviours to test by maximizing an information acquisition function that balances exploration (selecting points whose performance is uncertain) and exploitation (selecting points whose performance is expected to be high) (see Supplementary Methods). The selected behaviour is tested on the physical robot and the actual performance is recorded. The algorithm updates the expected performance of the tested behaviour and lowers its uncertainty. These updates are propagated to neighbouring solutions in the behavioural space by updating the Gaussian process (Supplementary Methods). These updated performance and confidence distributions affect which behaviour is tested next. This select–test–update loop repeats until the robot finds a behaviour whose measured performance is greater than some user-defined percentage (here, 90%) of the best performance predicted for any behaviour in the behaviour–performance map (Supplementary Methods).

We first tested our algorithm on a hexapod robot that needs to walk as fast as possible (Fig. 1b, d). The robot has 18 motors, an onboard computer, and a depth camera that allows the robot to estimate its walking speed (Supplementary Methods). The gait is parameterized by 36 real-valued parameters (Supplementary Methods) that describe, for each joint, the amplitude of oscillation, phase shift, and duty cycle (the fraction of the period that the joint angle is on one side of the midline).

(broken leg). **d**, After damage occurs, in this case making the robot unable to walk fast and in a straight line, damage recovery via IT&E begins. The robot tests different types of behaviours from an automatically generated map of the behaviour–performance space. After each test, the robot updates its predictions of which behaviours will perform well despite the damage. In this way, the robot rapidly discovers an effective compensatory behaviour.

The behaviour space is six-dimensional, where each dimension is the proportion of time the i th leg spends in contact with the ground (that is, the duty factor)¹ (Supplementary Methods).

The behaviour–performance map created contains approximately 13,000 different gaits (Supplementary Video 2 shows examples). We tested our robot under six different conditions: undamaged (C1 in Fig. 3a), four different structural failures (C2–C5 in Fig. 3a), and a temporary leg repair (C6 in Fig. 3a). We compare the walking speed of resultant gaits with a widely used, classic, hand-designed tripod gait¹ (Supplementary Methods). For each of the six damage conditions, we ran our adaptation step five times for each of eight independently generated behaviour–performance maps (with the default ‘duty factor’ behavioural description), leading to $6 \times 5 \times 8 = 240$ experiments in total. We also ran our adaptation step five times on eight independently generated behaviour–performance maps defined by an alternative behavioural description (‘body orientation’, see Supplementary Methods) on two damage conditions (Fig. 3b, c), leading to $2 \times 5 \times 8 = 80$ additional experiments.

When the robot is undamaged (C1 in Fig. 3a), our approach yields dynamic gaits that are 30% faster than the classic reference gait (Fig. 3b, median 0.32 m s^{-1} , 5th and 95th percentiles [0.26 m s^{-1} ; 0.36 m s^{-1}] versus 0.24 m s^{-1}), suggesting that IT&E is a good search algorithm for automatically producing successful robot behaviours, putting aside damage recovery. In all the damage scenarios, the reference gait is no longer effective (around 0.04 m s^{-1} for the four damage conditions, C2–C5 in Fig. 3b). After IT&E, the compensatory gaits achieve a reasonably fast speed ($>0.15 \text{ m s}^{-1}$) and are between three and seven times more efficient than the reference gait for that damage condition (gaits given in metres per second: 0.24 [0.18; 0.31] versus 0.04 for C2; 0.22 [0.18; 0.26] versus 0.03 for C3; 0.21 [0.17; 0.26] versus 0.04 for C4; 0.17 [0.12; 0.24] versus 0.05 for C5; and 0.3 [0.21; 0.33] versus 0.12 for C6).

These experiments demonstrate that IT&E allows the robot both to initially learn fast gaits and to recover reliably after physical damage.

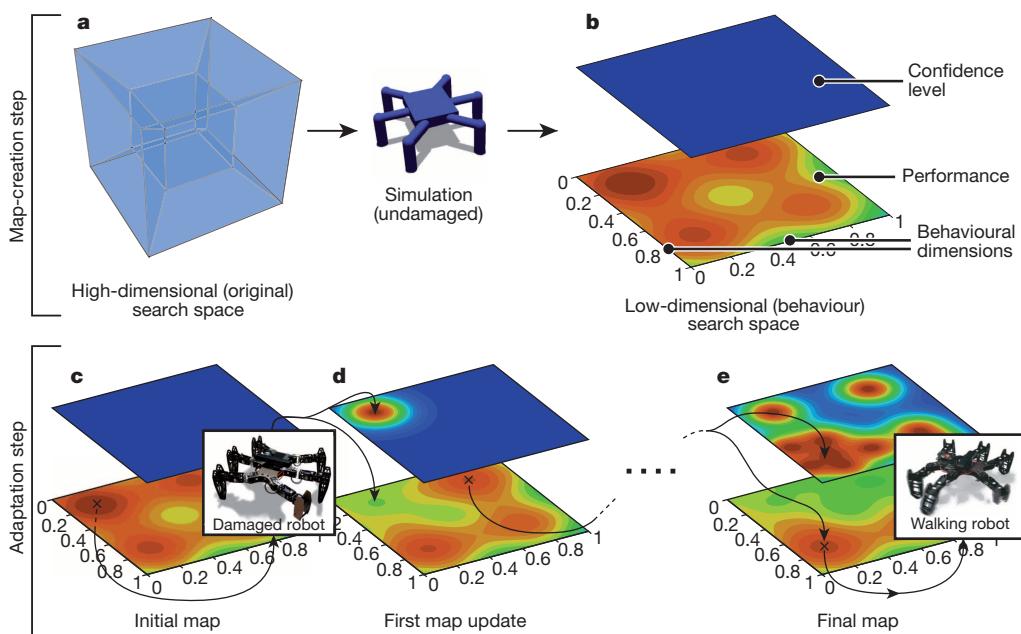


Figure 2 | The two steps of IT&E. **a, b,** Creating the behaviour–performance map. A user reduces a high-dimensional search space to a low-dimensional behaviour space by defining dimensions along which behaviours vary. In simulation, the high-dimensional space is then automatically searched to find a high-performing behaviour at each point in the low-dimensional behaviour space, creating a behaviour–performance map of the performance potential of each location in the low-dimensional space. In our hexapod robot experiments, the behaviour space is six-dimensional, with each dimension representing the portion of time that one leg is in contact with the ground. The confidence regarding the accuracy of the predicted performance for each behaviour in the behaviour–performance map is initially low because no tests on the physical robot have been conducted. **c–e,** Adaptation step: after damage, the robot selects a promising behaviour, tests it, updates the predicted performance of

that behaviour in the behaviour–performance map, and allocates a high confidence to this performance prediction. The predicted performances of nearby behaviours—and confidence in those predictions—are likely to be similar to the tested behaviour and are thus updated accordingly. This select–test–update loop is repeated until a tested behaviour on the physical robot performs better than 90% of the best predicted performance in the behaviour–performance map, a value that can decrease with each test (Extended Data Fig. 1). The algorithm that selects which behaviour to test next balances between choosing the behaviour with the highest predicted performance and behaviours that are different from those tested so far. Overall, the IT&E approach presented here rapidly locates which types of behaviours are least affected by the damage to find an effective, compensatory behaviour.

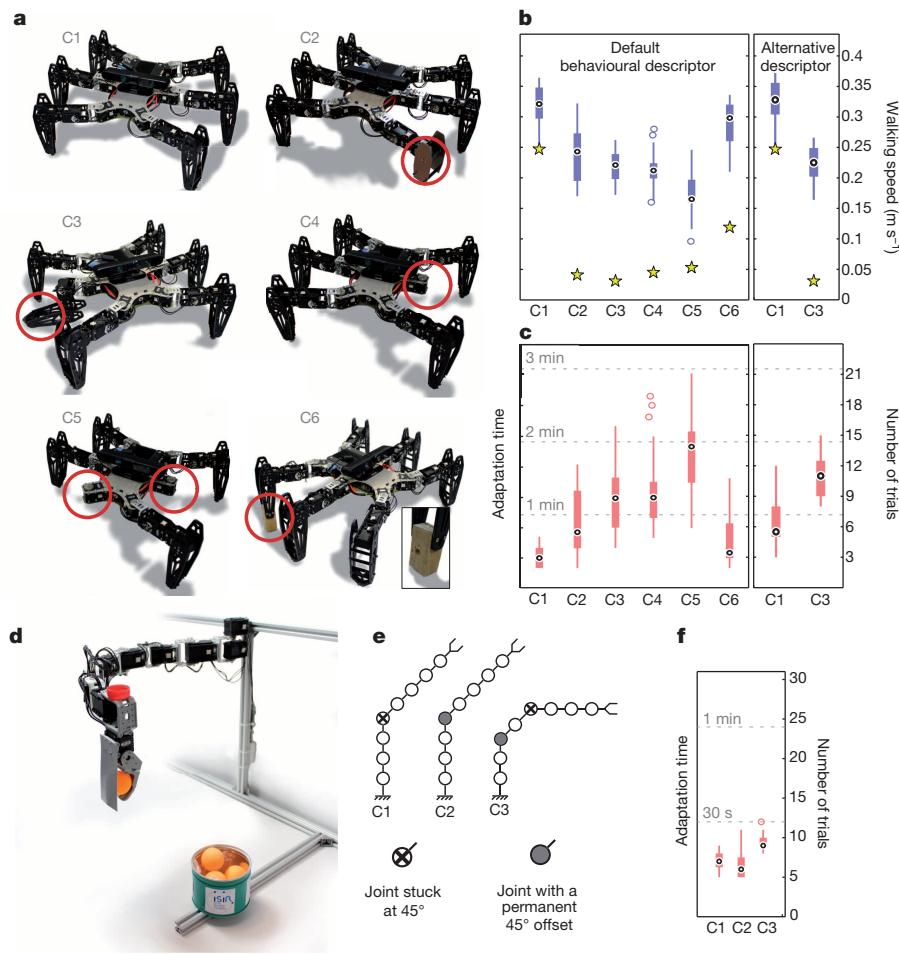
Additional experiments reveal that these capabilities are substantially faster than state-of-the-art algorithms (Bayesian optimization and policy gradient, Extended Data Fig. 2), and that IT&E can help with another major challenge in robotics: adapting to new environments (such as differently sloped terrain, Extended Data Fig. 3). On the undamaged or repaired robot (C6 in Fig. 3a), IT&E learns a walking gait in less than 30 s (Fig. 3c, the undamaged robot takes 24 s [16 s; 41 s] in 3 [2; 5] physical trials and the repaired robot takes 29 s [16 s; 82 s] in 3.5 [2; 10] trials). For the five damage scenarios, the robot adapts in approximately one minute (66 s [24 s; 134 s] in 8 [3; 16] trials). It is possible that for certain types of damage the prior information from the undamaged robot does not help, and could even hurt, in learning a compensatory behaviour (for example, if the map does not contain a compensatory behaviour). We did not find such a case in our experiments with the hexapod robot, but we did find a case on another robot in which the prior information provided little benefit (experiment 1 in the Supplementary Information).

Our results are qualitatively unchanged when using different behavioural characterizations, including randomly choosing six descriptors among 63 possibilities (Fig. 3b, c and Extended Data Fig. 4). Additional experiments show that reducing the high-dimensional parameter space to a low-dimensional behaviour space via the behaviour–performance map is the key component for IT&E: standard Bayesian optimization in the original parameter space does not find working controllers (gaits) (Extended Data Fig. 2). We investigated how the behaviour–performance map is updated when the robot loses a leg (C4 in Fig. 3a). Initially, the map predicts large areas of high performance. During adaptation, these areas disappear because the behaviours do not work well on the damaged robot. IT&E quickly identifies one of the

few remaining high-performance behaviours (Fig. 4 and Extended Data Figs 5 and 6).

The same damage recovery approach can be applied to any robot, such as a robotic arm. We tested 14 different damage conditions with a planar, eight-jointed robotic arm (Fig. 3d–f and Extended Data Fig. 7). The behaviour–performance map’s behavioural dimensions are the x , y position of the hand at the end of the arm. To show that the map-generating performance measure can be different from the ultimate performance measure, and to encourage smooth arm movements, the performance measure during map-generation is minimizing the variance of the eight specified motor angles (Supplementary Methods). During adaptation, performance is measured as distance to the target. As with the hexapod robot, our approach discovers a compensatory behaviour in less than 2 min, usually in less than 30 s, and with fewer than ten trials (Fig. 3f and Extended Data Fig. 7).

Although animals do not use the specific algorithm we present, there are parallels between IT&E and animal learning. Damage recovery in animals may occur without learning—for instance, due to the built-in robustness of evolved control loops²⁴—but if such pre-programmed robustness fails, many animals turn to learning¹⁹. As in our robot, such learning probably exploits an animal’s intuitions about how its intact body works to experiment with different behaviours to find what works best. Also like animals²⁵, IT&E allows the quick identification of working behaviours with a few, diverse tests instead of trying behaviours at random or trying small modifications to the best behaviour found so far. Additionally, the Bayesian optimization procedure followed by our robot appears similar to the technique employed by humans when they optimize an unknown function²³, and there is strong evidence that animal brains learn probability

**Figure 3 | Main experiments and results.**

a, Conditions tested on the physical hexapod robot. C1, The undamaged robot. C2, One leg is shortened by half. C3, One leg is unpowered. C4, One leg is missing. C5, Two legs are missing. C6, An imperfect, makeshift repair to the tip of one leg (by a human operator). **b**, Performance after adaptation. Box plots represent IT&E. The central circle is the median, the edges of the box are the 25th and 75th percentiles, the whiskers extend to the most extreme data points that are not considered outliers, and outliers are plotted individually. Yellow stars represent the performance of the handmade reference tripod gait (Supplementary Methods). Conditions C1–C6 are tested five times each for eight independently created behaviour–performance maps with the ‘duty factor’ behaviour description (that is, 40 experiments per damage condition, Supplementary Methods). Damage conditions C1 and C3 are also tested five times each for eight independently created behaviour–performance maps with the ‘body orientation’ behaviour description (Supplementary Methods). **c**, Time and number of trials required to adapt. Box plots represent IT&E. **d**, Robotic arm experiment. The eight-jointed, planar robot arm has to drop a ball into a bin. **e**, Example conditions tested on the physical robotic arm. C1, One joint is stuck at 45 degrees. C2, One joint has a permanent 45° offset. C3, One broken and one offset joint. Each of these damage conditions was replicated with 15 independent maps. A total of 14 conditions were tested (Extended Data Fig. 7). **f**, Time and number of trials required to reach within 5 cm of the bin centre. Each condition is tested with 15 independently created behaviour–performance maps.

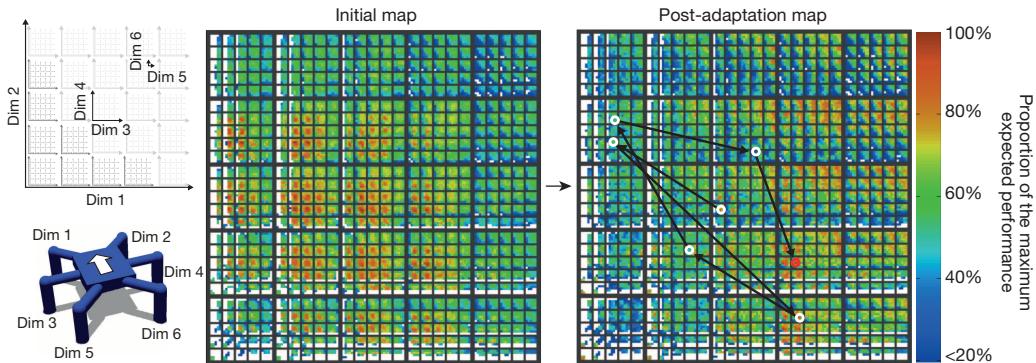


Figure 4 | An example behaviour–performance map. This map stores high-performing behaviours at each point in a six-dimensional behaviour space. Each dimension is the portion of time that each leg is in contact with the ground. The behavioural space is discretized at five values for each dimension (0; 0.25; 0.5; 0.75 and 1). Each coloured pixel represents the highest-performing behaviour discovered during map creation at that point in the behaviour space. The matrices visualize the six-dimensional behavioural space in two dimensions according to the legend at the top left. The behaviour–performance map is created with a simulated robot (bottom left) in the Open Dynamics Engine physics simulator (<http://www.ode.org>). The left matrix is a

pre-adaptation map produced by the map creation algorithm. During adaptation, the map is updated as tests are conducted (in this case, for the damage condition where the robot is missing one leg: C4 in Fig. 3a). The right matrix shows the state of the map after a compensatory behaviour is discovered. The arrows and white circles represent the order in which behaviours were tested on the physical robot. The red circle is the final, discovered, compensatory behaviour. Among other areas, high-performing behaviours can be found for the damaged robot in the first two columns of the third dimension. These columns represent behaviours that use the central-left leg least, which is the leg that is missing.

distributions, combine them with prior knowledge, and act as Bayesian optimizers^{26,27}.

An additional parallel is that IT&E primes the robot for creativity during a motionless period, after which the generated ideas are tested.

This process is reminiscent of the finding that some animals start the day with new ideas that they may quickly disregard after experimenting with them²⁸, and more generally, that sleep improves creativity on cognitive tasks²⁹. A final parallel is that the simulator and Gaussian

process components of IT&E are two forms of predictive models, which are known to exist in animals^{14,30}. Overall, we have shown that IT&E has parallels in biology and makes robots behave more like animals by endowing them with the ability to adapt rapidly to unforeseen circumstances.

Received 24 December 2014; accepted 17 March 2015.

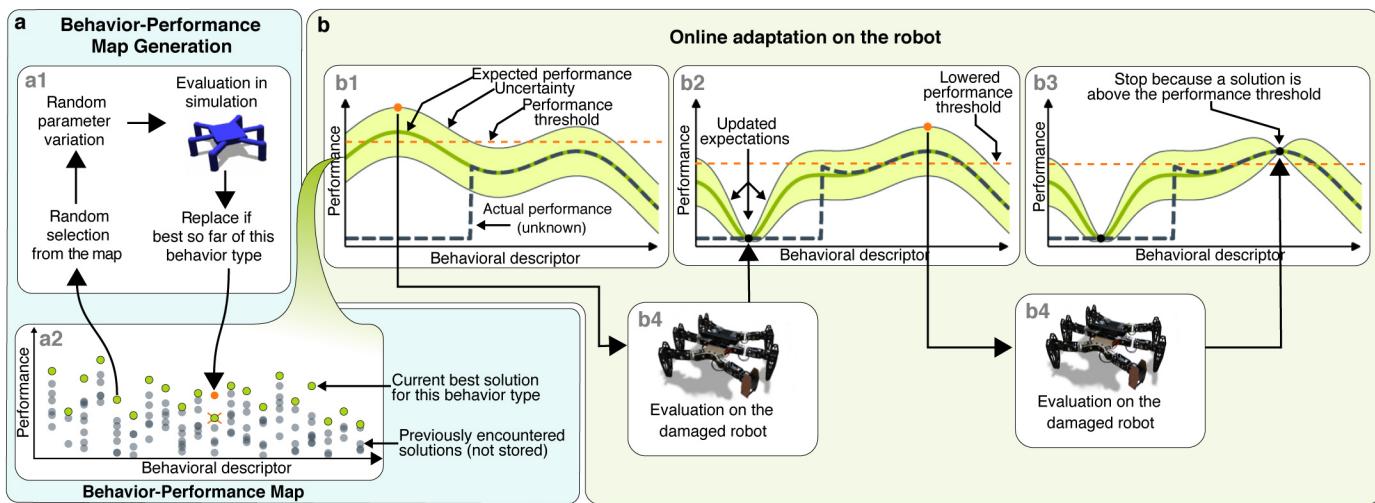
1. Siciliano, B. & Khatib, O. *Springer Handbook of Robotics* (Springer, 2008).
2. Murphy, R. R. Trial by fire. *Robot. Automat. Mag.* **11**, 50–61 (2004).
3. Nagatani, K. *et al.* Emergency response to the nuclear accident at the Fukushima Daiichi nuclear power plants using mobile rescue robots. *J. Field Robot.* **30**, 44–63 (2013).
4. Broadbent, E., Stafford, R. & MacDonald, B. Acceptance of healthcare robots for the older population: review and future directions. *Int. J. Social Robot.* **1**, 319–330 (2009).
5. Thrun, S. *et al.* Stanley: the robot that won the DARPA grand challenge. *J. Field Robot.* **23**, 661–692 (2006).
6. Sanderson, K. Mars rover Spirit (2003–10). *Nature* **463**, 600 (2010).
7. Antonelli, G., Fossen, T. I. & Yoerger, D. R. In *Springer Handbook of Robotics* (eds Siciliano, B. & Khatib, O.) 987–1008 (Springer, 2008).
8. Carlson, J. & Murphy, R. R. How UGVs physically fail in the field. *IEEE Trans. Robot.* **21**, 423–437 (2005).
9. Blanke, M., Kinnaert, M., Lunze, J. & Staroswiecki, M. *Diagnosis and Fault-Tolerant Control* (Springer, 2006).
10. Sproewitz, A., Moeckel, R., Maye, J. & Ijspeert, A. Learning to move in modular robots using central pattern generators and online optimization. *Int. J. Robot. Res.* **27**, 423–443 (2008).
11. Christensen, D. J., Schultz, U. P. & Stoy, K. A distributed and morphology-independent strategy for adaptive locomotion in self-reconfigurable modular robots. *Robot. Auton. Syst.* **61**, 1021–1035 (2013).
12. Kober, J., Bagnell, J. A. & Peters, J. Reinforcement learning in robotics: a survey. *Int. J. Robot. Res.* **32**, 1238–1274 (2013).
13. Verma, V., Gordon, G., Simmons, R. & Thrun, S. Real-time fault diagnosis. *Robot. Automat. Mag.* **11**, 56–66 (2004).
14. Bongard, J., Zykov, V. & Lipson, H. Resilient machines through continuous self-modeling. *Science* **314**, 1118–1121 (2006).
15. Kluger, J. & Lovell, J. *Apollo 13* (Mariner Books, 2006).
16. Jarvis, S. L. *et al.* Kinematic and kinetic analysis of dogs during trotting after amputation of a thoracic limb. *Am. J. Vet. Res.* **74**, 1155–1163 (2013).
17. Fuchs, A., Goldner, B., Nolte, I. & Schilling, N. Ground reaction force adaptations to tripodal locomotion in dogs. *Vet. J.* **201**, 307–315 (2014).
18. Argall, B. D., Chernova, S., Veloso, M. & Browning, B. A survey of robot learning from demonstration. *Robot. Auton. Syst.* **57**, 469–483 (2009).
19. Wolpert, D. M., Ghahramani, Z. & Flanagan, J. R. Perspective and problems in motor learning. *Trends Cogn. Sci.* **5**, 487–494 (2001).
20. Santello, M. Postural hand synergies for tool use. *J. Neurosci.* **18**, 10105–10115 (1998).
21. Rasmussen, C. E. & Williams, C. K. I. *Gaussian Processes for Machine Learning* (MIT Press, 2006).
22. Mockus, J. *Bayesian Approach to Global Optimization: Theory and Applications* (Kluwer Academic, 2013).
23. Borji, A. & Itti, L. Bayesian optimization explains human active search. *Adv. Neural Inform. Process. Syst.* **26**, 55–63 (2013).
24. Grillner, S. The motor infrastructure: from ion channels to neuronal networks. *Nature Rev. Neurosci.* **4**, 573–586 (2003).
25. Benson-Amram, S. & Holekamp, K. E. Innovative problem solving by wild spotted hyenas. *Proc. R. Soc. Lond. B* **279**, 4087–4095 (2012).
26. Pouget, A., Beck, J. M., Ma, W. J. & Latham, P. E. Probabilistic brains: knowns and unknowns. *Nature Neurosci.* **16**, 1170–1178 (2013).
27. Kording, K. P. & Wolpert, D. M. Bayesian integration in sensorimotor learning. *Nature* **427**, 244–247 (2004).
28. Derégnaucourt, S., Mitra, P. P., Fehér, O., Pytte, C. & Tchernichovski, O. How sleep affects the developmental learning of bird song. *Nature* **433**, 710–716 (2005).
29. Wagner, U., Gais, S., Haider, H., Verleger, R. & Born, J. Sleep inspires insight. *Nature* **427**, 352–355 (2004).
30. Ito, M. Control of mental activities by internal models in the cerebellum. *Nature Rev. Neurosci.* **9**, 304–313 (2008).

Supplementary Information is available in the online version of the paper.

Acknowledgements We thank L. Tedesco, S. Doncieux, N. Bredeche, S. Whiteson, R. Calandra, J. Droulez, P. Bessière, F. Lesaint, C. Thurat, S. Ivaldi, C. Lan Sun Luk, J. Li, J. Huizinga, R. Velez, H. Mengistu, M. Norouzzadeh, T. Clune, and A. Nguyen for feedback and discussions. This work has been funded by the ANR Creadapt project (ANR-12-JS03-0009), the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement number 637972), and a Direction Générale de l'Armement (DGA) scholarship to A.C.

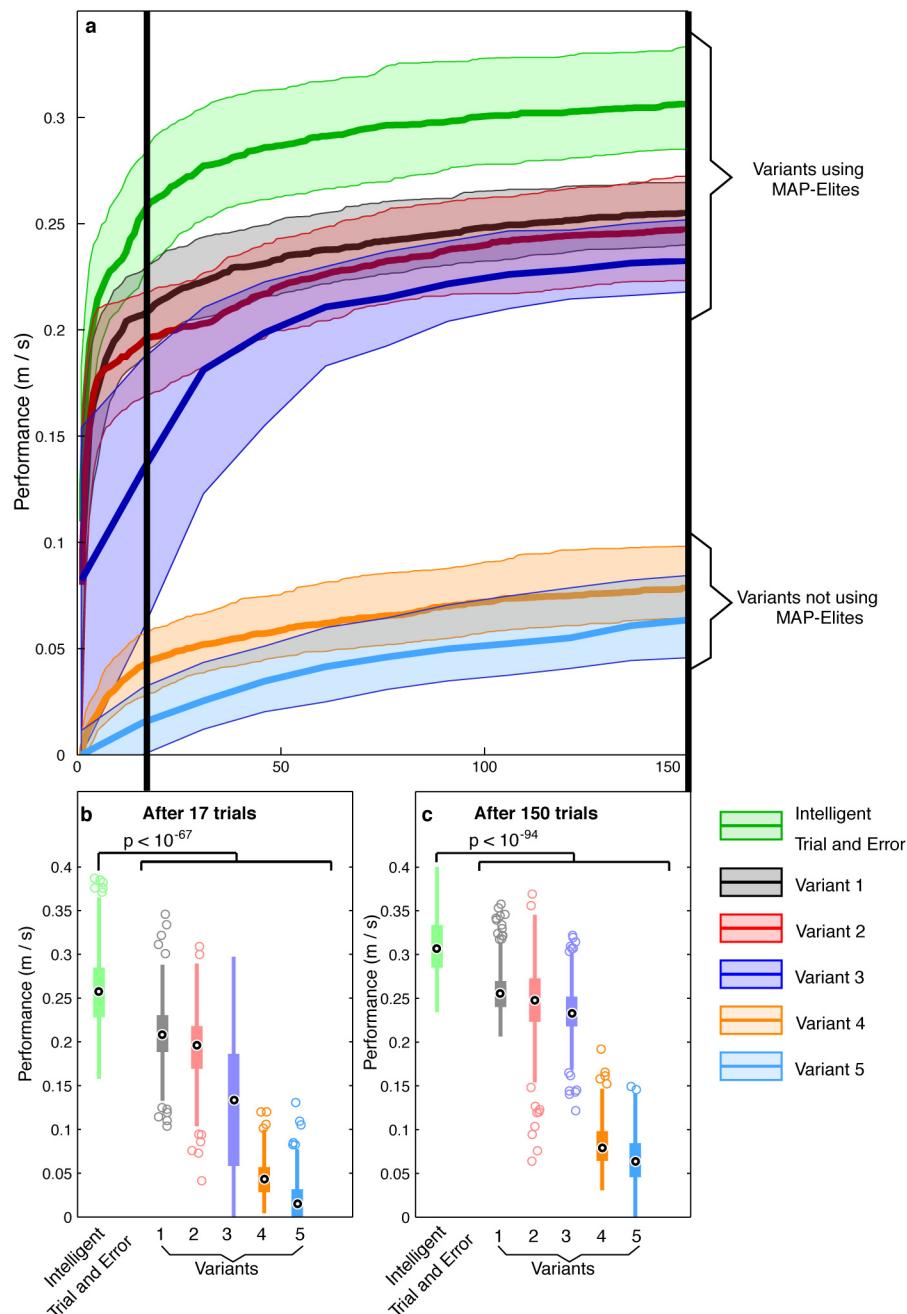
Author Contributions A.C. and J.-B.M. designed the study. A.C. and D.T. performed the experiments. A.C., J.-B.M., D.T. and J.C. analysed the results, discussed additional experiments, and wrote the paper.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to J.-B.M. (jean-baptiste.mouret@inria.fr).



Extended Data Figure 1 | An overview of the IT&E algorithm. **a**, Behaviour-performance map creation. After being initialized with random controllers, the behavioural map (a2), which stores the highest-performing controller found so far of each behaviour type, is improved by repeating the process depicted in a1 until newly generated controllers are rarely good enough to be added to the map (here, after 40 million evaluations). This step, which occurs in simulation (and required roughly two weeks on one multi-core computer; see Supplementary Methods, section ‘Running time’), is computationally expensive, but only needs to be performed once per robot (or robot design) before the robot is deployed. **b**, Adaptation. In b1, each behaviour from the behaviour–performance map has an expected performance based on its performance in simulation (dark green line) and an estimate of uncertainty regarding this predicted performance (light green band). The actual performance on the now-damaged robot (black dashed line) is unknown to the algorithm. A behaviour is selected for the damaged robot to try. This selection is made by balancing exploitation—trying

behaviours expected to perform well—and exploration—trying behaviours whose performance is uncertain (Supplementary Methods, section ‘Acquisition function’). Because all points initially have equal, maximal uncertainty, the first point chosen is that with the highest expected performance. Once this behaviour is tested on the physical robot (b4), the performance predicted for that behaviour is set to its actual performance, the uncertainty regarding that prediction is lowered, and the predictions for, and uncertainties about, nearby controllers are also updated (according to a Gaussian process model, see Supplementary Methods, section ‘Kernel function’), the results of which can be seen in b2. The process is then repeated until performance on the damaged robot is 90% or greater of the maximum expected performance for any behaviour (b3). This performance threshold (orange dashed line) lowers as the maximum expected performance (the highest point on the dark green line) is lowered, which occurs when physical tests on the robot underperform expectations, as occurred in b2.

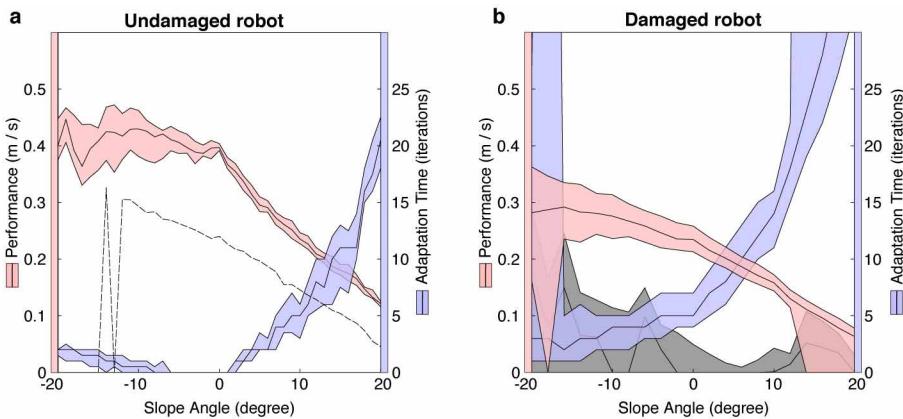


Variant	Behavior-performance map creation	Priors on performance	Search algorithm	equivalent approach
Intelligent Trial and Error	MAP-Elites	yes	Bayesian Optimization	-
Variant 1	MAP-Elites	none	Random Search	-
Variant 2	MAP-Elites	none	Bayesian Optimization	-
Variant 3	MAP-Elites	none	Policy Gradient	-
Variant 4	none	none	Bayesian Optimization	Lizotte et al. (2007)
Variant 5	none	none	Policy Gradient	Kohl et al. (2004)

Extended Data Figure 2 | The contribution of each subcomponent of the IT&E algorithm. **a**, Adaptation progress versus the number of robot trials. The walking speed achieved with IT&E and several ‘knockout’ variants that are missing one of the algorithm’s key components (defined in the table below the plots). Some variants (4 and 5) correspond to state-of-the-art learning algorithms (policy gradient³¹ and Bayesian optimization^{32,33,34}). MAP-Elites (written by J.-B.M. and J.C.; ref. 35) is the algorithm that IT&E uses to create the behaviour–performance map (see Supplementary Methods). The bold lines represent the medians and the coloured areas extend to the 25th and 75th percentiles. **b, c**, Adaptation performance after 17 and 150 trials. Shown is the

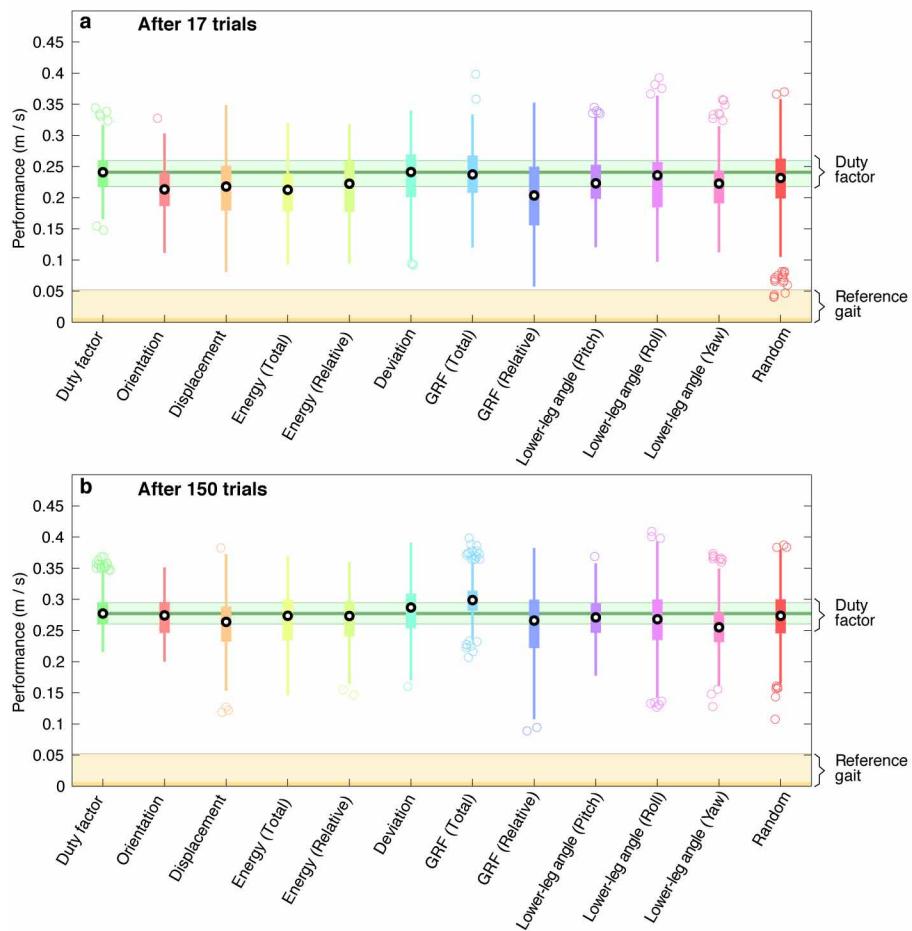
speed of the compensatory behaviour discovered by each algorithm after 17 and 150 evaluations on the simulated robot, respectively. For all panels, data are pooled across six damage conditions (the removal of each of the six legs in turn) for eight independently generated maps and replicated ten times (for a total of 480 replicates). *P* values are computed via a Wilcoxon ranksum test. The central circle is the median, the edges of the box show the 25th and 75th percentiles, the whiskers extend to the most extreme data points that are not considered outliers, and outliers are plotted individually. See experiment 2 in the Supplementary Information for methods and analysis.

31. Kohl, N. & Stone, P. Policy gradient reinforcement learning for fast quadrupedal locomotion. In *Proc. IEEE Int. Conf. on ‘Robotics and Automation’ (ICRA)* 2619–2624 (IEEE, 2004).
32. Lizotte, D. J., Wang, T., Bowling, M. H. & Schuurmans, D. Automatic gait optimization with Gaussian process regression. In *Proc. Int. Joint Conf. on ‘Artificial Intelligence’ (IJCAI)* 944–949 (2007).
33. Tesch, M., Schneider, J. & Choset, H. Using response surfaces and expected improvement to optimize snake robot gait parameters. In *Proc. IEEE/RSJ Int. Conf. on ‘Intelligent Robots and Systems (IROS)’* 1069–1074 (IEEE, 2011).
34. Calandra, R., Seyfarth, A., Peters, J. & Deisenroth, M. P. An experimental comparison of bayesian optimization for bipedal locomotion. In *Proc. IEEE Int. Conf. on ‘Robotics and Automation’ (ICRA)* 1951–1958 (IEEE, 2014).
35. Mouret, J.-B. & Clune, J. Illuminating search spaces by mapping elites. Preprint at <http://arxiv.org/abs/1504.04909> (2015).



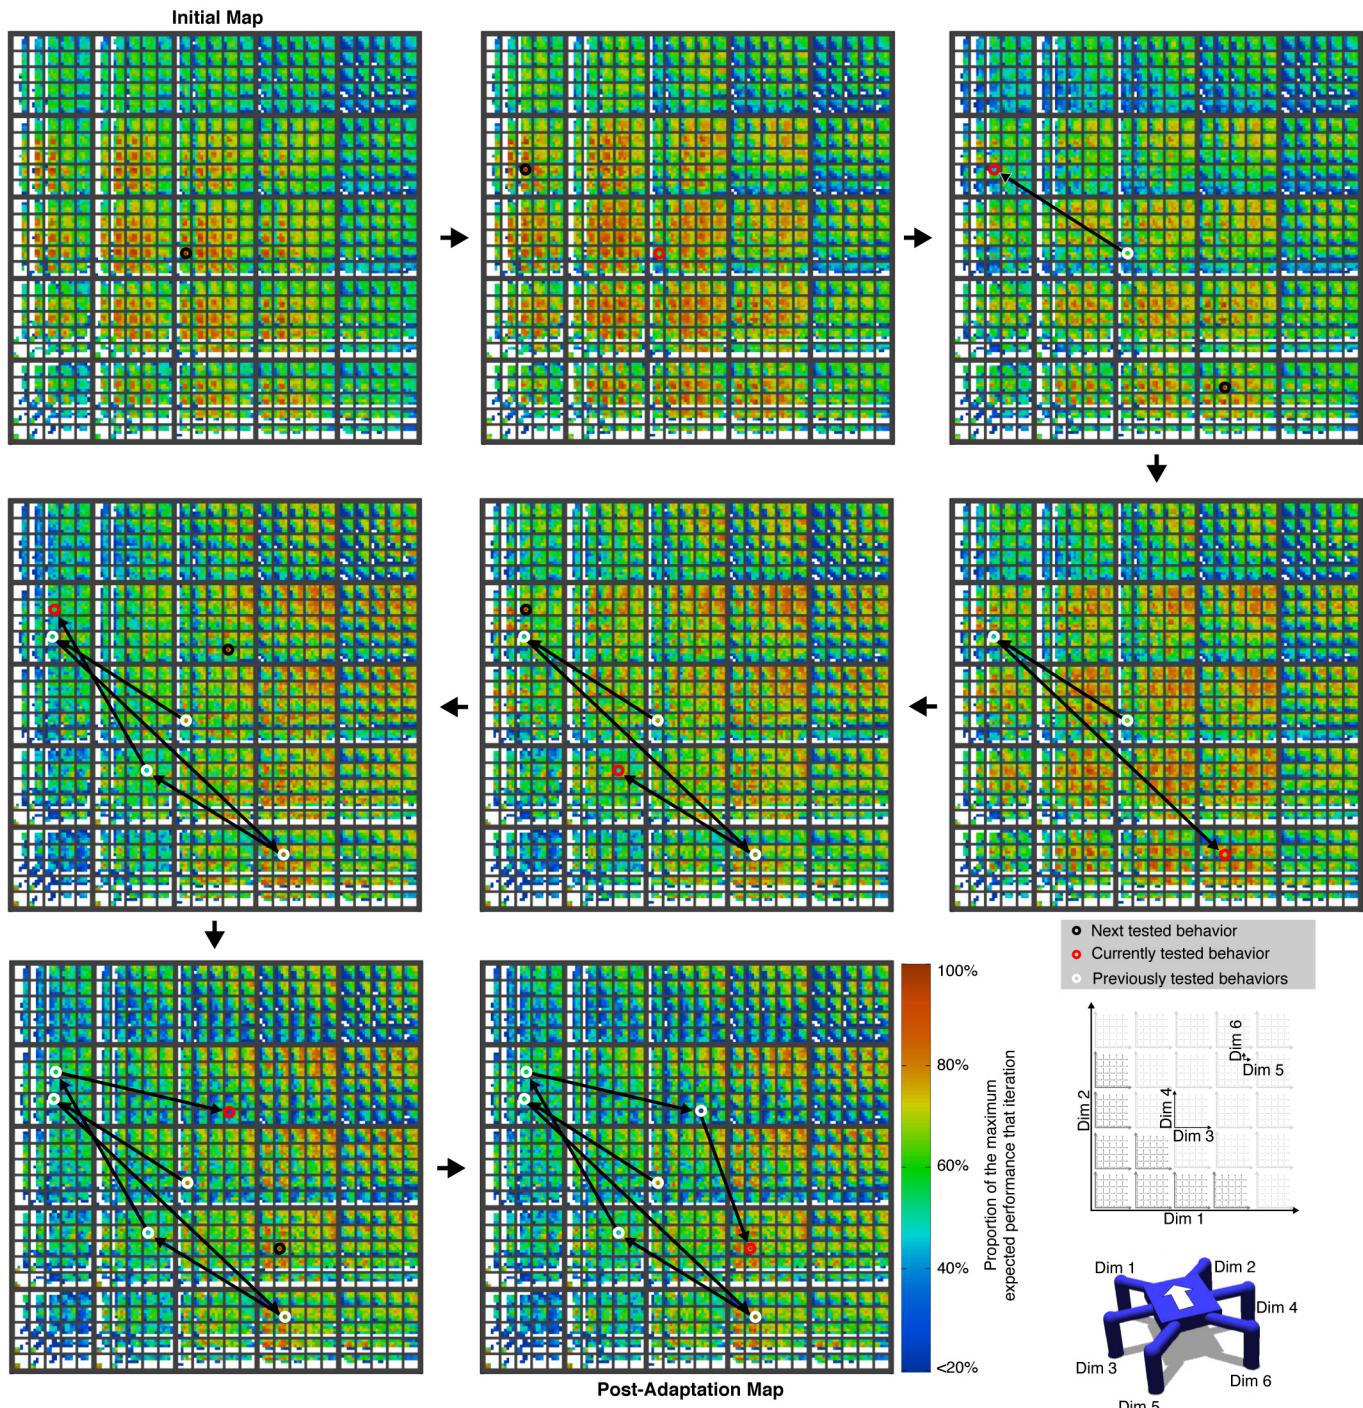
Extended Data Figure 3 | The IT&E algorithm is robust to environmental changes. Each plot shows both the performance and required adaptation time for IT&E when the robot must adapt to walk on terrains of different slopes. **a**, Adaptation performance on an undamaged, simulated robot. On all slope angles, with very few trials, the IT&E algorithm (pink shaded region) finds fast gaits that outperform the reference gait (black dashed line). We performed ten replicates for each of the eight maps and each one-degree increment between -20° and $+20^\circ$ degrees (a total of 3,280 experiments). **b**, Adaptation performance on a damaged, simulated robot. The robot is damaged by having each of its six legs removed in six different damage scenarios. Data

are pooled from all six of these damage conditions. For each damage condition, we performed ten replicates for each of the eight maps and each two-degree increment between -20° and $+20^\circ$ degrees (a total of 10,080 experiments). The median compensatory behaviour found via IT&E outperforms the median reference controller on all slope angles. The middle, black lines represent medians, while the colored areas extend to the 25th and 75th percentiles. In **a**, the black dashed line is the performance of a classic tripod gait for reference. In **b**, the reference gait is tried in all six damage conditions and its median (black line) and 25th and 75th percentiles (black coloured area) are shown. See experiment 3 in the Supplementary Information for methods and analysis.



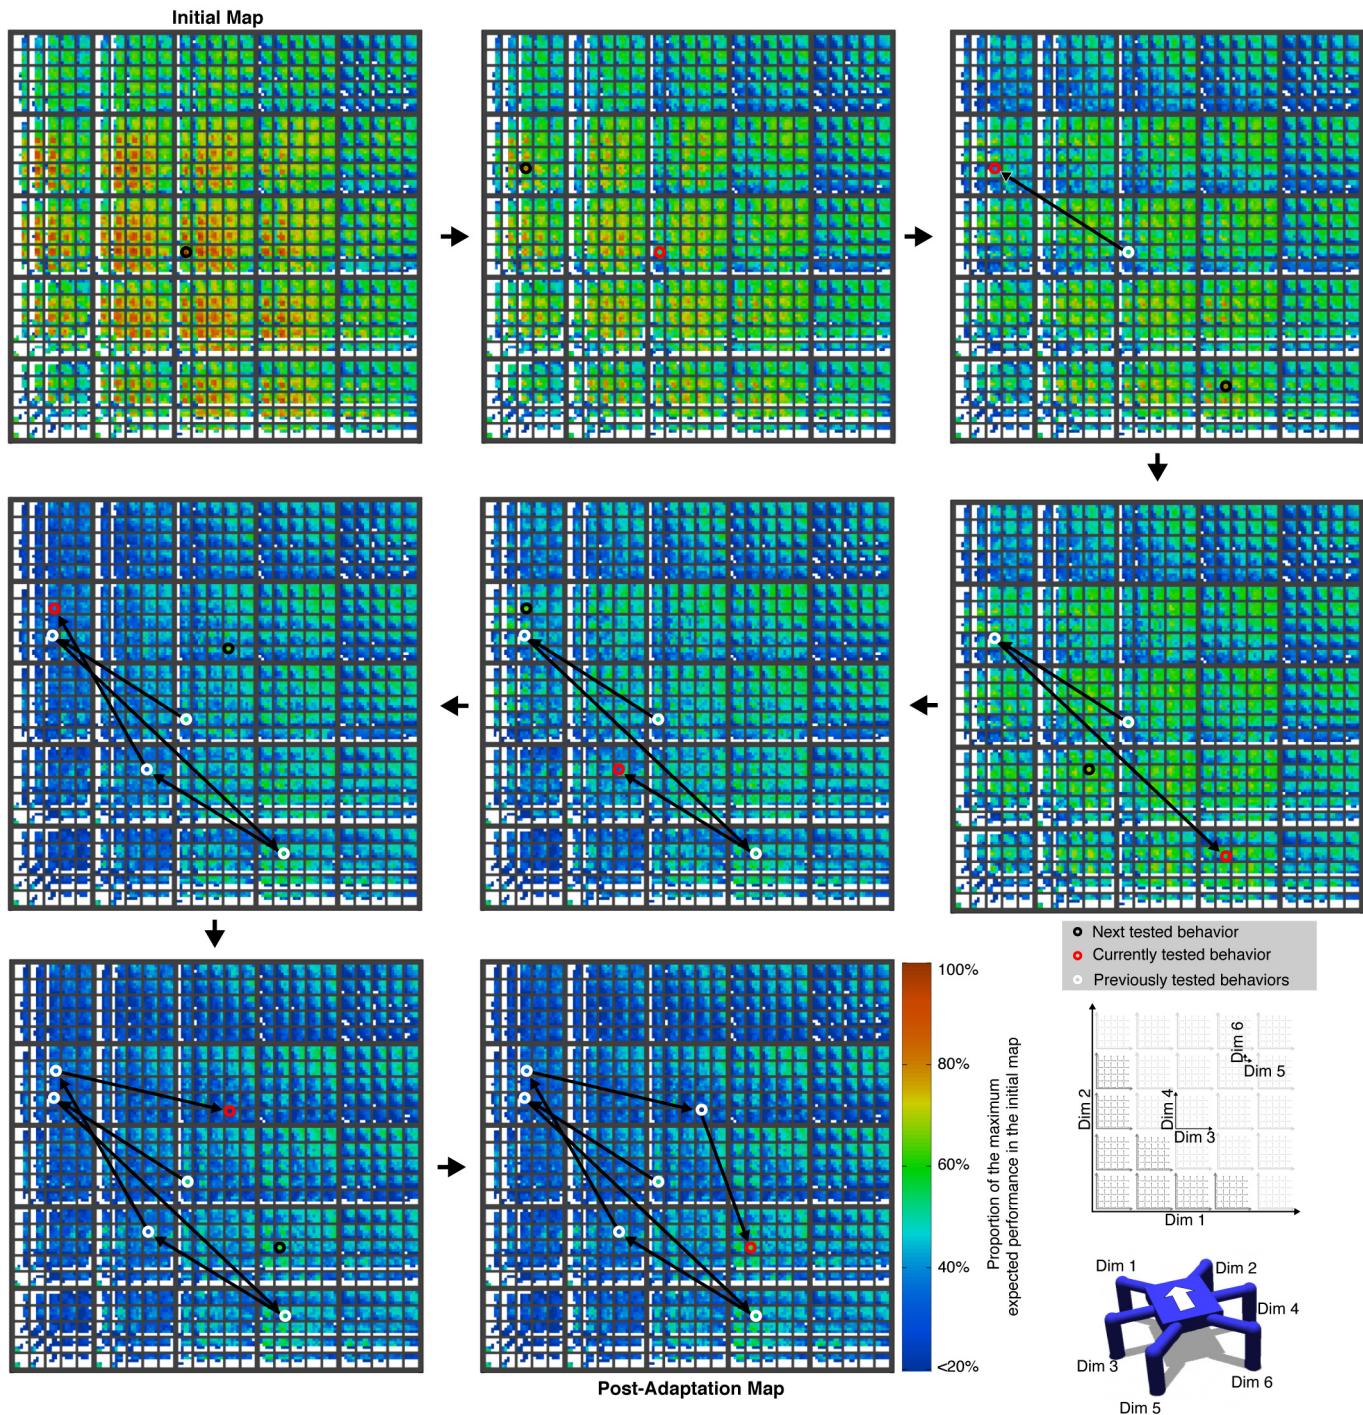
Extended Data Figure 4 | The IT&E algorithm is largely robust to alternative choices of behaviour descriptors. **a, b,** The speed of the compensatory behaviour discovered by IT&E for various choices of behaviour descriptors. Performance is plotted after 17 and 150 evaluations in **a** and **b**, respectively. Experiments were performed on a simulated, damaged hexapod. The damaged robot has each of its six legs removed in six different damage scenarios. Data are pooled across all six damage conditions. As described in experiment 5 of the Supplementary Information, the evaluated behaviour descriptors characterize the following: (1) time each leg is in contact with the ground ('Duty factor'); (2) orientation of the robot frame ('Orientation'); (3) instantaneous velocity of the robot ('Displacement'); (4) energy expended by the robot in walking ('Energy (Total)', 'Energy (Relative)'); (5) deviation from a straight line ('Deviation'); (6) ground reaction force on each leg ('GRF (Total)', 'GRF (Relative)'); (7) the angle of each leg when it touches the ground ('Lower-leg angle (Pitch)', 'Lower-leg angle (Roll)', 'Lower-leg angle (Yaw)'); and (8) a random selection without replacement from subcomponents of all the available behaviour descriptors (1)–(7) ('Random').

We created eight independently generated maps for each of 11 intentionally chosen behavioral descriptors (88 in total). We created one independently generated map for each of 20 randomly chosen behavioural descriptors (20 in total). For each of these 108 maps, we performed ten replicates of IT&E for each of the six damage conditions. In total, there were thus $8 \times 11 \times 10 \times 6 = 5,280$ experiments for the intentionally chosen behavioural descriptors and $1 \times 20 \times 10 \times 6 = 1,200$ experiments for the randomly chosen behavioural descriptors. For the hand-designed reference gait (yellow) and the compensatory gaits found by the default duty factor behaviour descriptor (green), the bold lines represent the medians and the coloured areas extend to the 25th and 75th percentiles of the data. For the other treatments, including the duty factor treatment, black circles represent the median, the edges of the boxes show the 25th and 75th percentiles, the whiskers extend to the most extreme data points that are not considered outliers, and outliers are plotted individually. See experiment 5 in the Supplementary Information for methods and analysis.



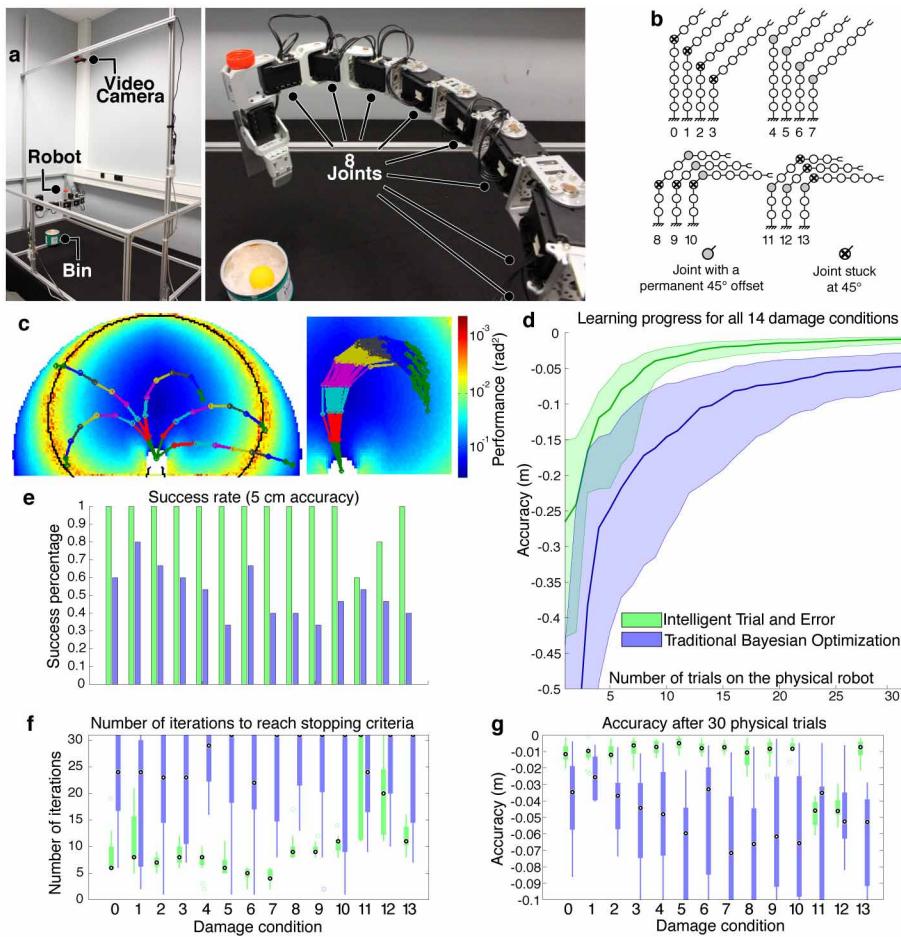
Extended Data Figure 5 | How the behaviour performance map (data normalized) is explored to discover a compensatory behaviour. Data are normalized each iteration to highlight the range of remaining performance predictions. The robot is adapting to damage condition C4 (see Fig. 3a). Colours represent the performance prediction for each point in the map relative to the highest performing prediction in the map at that step of the process. A

black circle indicates the next behaviour to be tested on the physical robot. A red circle indicates the behaviour that was just tested (note that the performance predictions surrounding it have changed versus the previous panel). Arrows reveal the order that points have been explored. The red circle in the last map is the final, selected, compensatory behaviour. In this scenario, the robot loses leg number 3. The six-dimensional space is visualized according to the inset legend.



Extended Data Figure 6 | How the behaviour performance map (data not normalized) is explored to discover a compensatory behaviour. The data are not normalized in order to highlight that performance predictions decrease as it is discovered that predictions from the simulated, undamaged robot do not work well on the damaged robot. The robot is adapting to damage condition C4 (see Fig. 3a). Colours represent the performance prediction for each point in the map relative to the highest performing prediction in the first map. A black circle indicates the next behaviour to be tested on the physical robot. A red

circle indicates the behaviour that was just tested (note that the performance predictions surrounding it have changed versus the previous panel). Arrows reveal the order that points have been explored. The red circle in the last map in the sequence is the final, selected, compensatory behaviour. In this scenario, the robot loses leg number 3. The six-dimensional space is visualized according to the inset legend. The data visualized in this figure are identical to those in the previous figure: the difference is simply whether the data are renormalized for each new map in the sequence.



Extended Data Figure 7 | IT&E works on a completely different type of robot (the robotic arm experiment). **a**, The robotic arm experimental setup. **b**, Tested damage conditions. **c**, Example of behaviour performance maps (colour maps) and behaviours (overlaid arm configurations) obtained with MAP-Elites. Left: A typical behaviour–performance map produced by MAP-Elites with five example behaviours, where a behaviour is described by the angle of each of the eight joints. The colour of each point is a function of its performance (in radians squared), which is defined as having low variance in the joint angles (that is, a zigzag arm has lower performance than a straighter arm that reaches the same point). Right: Neighbouring points in the map tend to have similar behaviours, thanks to the performance function, which would penalize more jagged ways of reaching those points. That neighbours have similar behaviours justifies updating predictions about the performance of nearby behaviours after testing a single behaviour on the real (damaged) robot. **d**, Accuracy (in metres) versus trial number for IT&E and traditional

Bayesian optimization. The experiment was conducted on the physical robot, with 15 independent replications for each of the 14 damage conditions. Accuracy is pooled from all of these $14 \times 15 = 210$ experiments for each algorithm. The middle lines represent medians, while the coloured areas extend to the 25th and 75th percentiles. **e**, Success for each damage condition. Shown is the success rate for the 15 replications for each damage condition, defined as the percentage of replicates in which the robot reaches within 5 cm of the bin centre. **f**, Trials required to adapt. Shown is the number of iterations required to reach within 5 cm of the bin centre. **g**, Accuracy after 30 physical trials for each damage condition (with the stopping criterion disabled). For **f** and **g**, the central circle is the median, the edges of the box show the 25th and 75th percentiles, the whiskers extend to the most extreme data points that are not considered outliers, and outliers are plotted individually. See experiment 1 in the Supplementary Information for methods and analysis.

Ocean impact on decadal Atlantic climate variability revealed by sea-level observations

Gerard D. McCarthy¹, Ivan D. Haigh², Joël J.-M. Hirschi¹, Jeremy P. Grist¹ & David A. Smeed¹

Decadal variability is a notable feature of the Atlantic Ocean and the climate of the regions it influences. Prominently, this is manifested in the Atlantic Multidecadal Oscillation (AMO) in sea surface temperatures. Positive (negative) phases of the AMO coincide with warmer (colder) North Atlantic sea surface temperatures. The AMO is linked with decadal climate fluctuations, such as Indian and Sahel rainfall¹, European summer precipitation², Atlantic hurricanes³ and variations in global temperatures⁴. It is widely believed that ocean circulation drives the phase changes of the AMO by controlling ocean heat content⁵. However, there are no direct observations of ocean circulation of sufficient length to support this, leading to questions about whether the AMO is controlled from another source⁶. Here we provide observational evidence of the widely hypothesized link between ocean circulation and the AMO. We take a new approach, using sea level along the east coast of the United States to estimate ocean circulation on decadal timescales. We show that ocean circulation responds to the first mode of Atlantic atmospheric forcing, the North Atlantic Oscillation, through circulation changes between the subtropical and subpolar gyres—the intergyre region⁷. These circulation changes affect the decadal evolution of North Atlantic heat content and, consequently, the phases of the AMO. The Atlantic overturning circulation is declining⁸ and the AMO is moving to a negative phase. This may offer a brief respite from the persistent rise of global temperatures⁴, but in the coupled system we describe, there are compensating effects. In this case, the negative AMO is associated with a continued acceleration of sea-level rise along the northeast coast of the United States^{9,10}.

The difficulty in linking ocean circulation changes to decadal climate variations lies in the fact that long observational records of ocean transports are rare. Measurements such as those of the Florida Current since 1982¹¹ and the Greenland–Scotland ridge transports¹² since the mid-1990s are some of the longest continuous ocean transport records available. Continuous, full-depth, basin-wide measurements of the Atlantic overturning circulation only began in 2004 with the RAPID monitoring project at 26° N (ref. 13). None of these records are long enough to directly link ocean circulation with decadal climate variations such as the AMO.

Sea-level measurements from tide gauges provide an integrated measure of water column properties and offer timeseries of sufficient length (Extended Data Fig. 1) to study decadal ocean circulation variations. Investigating ocean circulation using tide gauges is not new: the first attempt to estimate the Gulf Stream using tide gauges was made in 1938¹⁴. The principle is based on geostrophic dynamics: on timescales longer than a few days, ocean circulation is in geostrophic balance so, looking downstream, the sea level is seen to increase from left to right in the Northern Hemisphere.

Estimates of the Gulf Stream using tide gauges have focused on the American east coast, with an offshore estimate of sea level from either an island gauge¹⁵ or a reconstructed sea level¹⁶. A weakness of this method is that the offshore measurement lies in the eddy-filled ocean

where sea-level fluctuations at any one point are influenced by mesoscale variations¹⁷ even on long timescales, increasing the difficulty of making estimates of ocean circulation that are coherent on large spatial scales. This is the case for sea level at Bermuda, whose decadal fluctuations can be reproduced by considering a Rossby wave response to wind forcing¹⁶. To make estimates of ocean circulation that capture the fluctuations in large-scale circulation and are less affected by eddy variability, measurements close to or on the western boundary are necessary¹⁸. We account for this by focusing on the gradient of sea level along the US east coast. Here the mean dynamic sea level decreases to the north (Fig. 1a) due to the transition from subtropical to subpolar gyres. This dynamic gradient reflects a circulation that contains elements not only of the Gulf Stream but also of cold, subpolar water from the north, primarily associated with the overturning circulation¹⁹. Indeed, in model simulations, this meridional gradient of sea level along this coast responds strongly to declines in the Atlantic overturning circulation²⁰. Ultimately, it is the heat transport that we are interested in. And while the overturning circulation carries about 90% of the heat at subtropical latitudes²¹, ocean heat transport at the latitude of the intergyre region consists of similar contributions from both the overturning circulation and the gyre circulation²². For this reason, we do not discuss separately overturning and gyre but only ocean circulation in this intergyre region, which contains elements of both mechanisms.

Sea-level fluctuations from Florida to Boston divide into two coherent groups either side of Cape Hatteras²³ (Extended Data Figs 2, 3). This large-scale coherence in sea level is driven by ocean circulation. North of Cape Hatteras, coherent sea-level fluctuations have been linked with fluctuations in the overturning circulation^{19,24}. South of Cape Hatteras, fluctuations in the Gulf Stream from Florida to Cape Hatteras are reflected in sea-level fluctuations. As Cape Hatteras marks the boundary between the subtropical and subpolar gyres on this coastline (Fig. 1a), we can construct a single sea-level composite representative of the subtropical (subpolar) circulation by averaging sea level from linearly detrended, deseasonalized tide gauges, with the inverse barometer effect removed, south (north) of the Cape (Fig. 1b, c). The difference, south minus north (Fig. 1d), represents our circulation index. This index projects onto observed surface velocities during the satellite era in the intergyre region, with a positive index associated with more northwards flow and a more northerly path of this circulation (Extended Data Fig. 4). Similarly, in a high-resolution ocean model, over timescales that contain both the cool phase of the AMO in the 1970s²⁵ and the warm phase of the 1990s²⁶, the sea-level index projects onto a similar pattern of circulation, with a positive index associated with more northward heat transport (Extended Data Fig. 5).

Ocean circulation is proportional to heat transport at both subtropical and subpolar latitudes²². A number of recent studies (see, for example, ref. 27) have emphasized the dominant role of ocean heat transport in heat content changes, relating the accumulation (in time) of heat transport to heat content. This suggests that the accumulation of our sea-level index across Cape Hatteras, as a proxy for ocean circulation, can be related to ocean heat content. The largest AMO

¹National Oceanography Centre, University of Southampton Waterfront Campus, European Way, Southampton SO14 3ZH, UK. ²Ocean and Earth Science, National Oceanography Centre, University of Southampton Waterfront Campus, European Way, Southampton SO14 3ZH, UK.

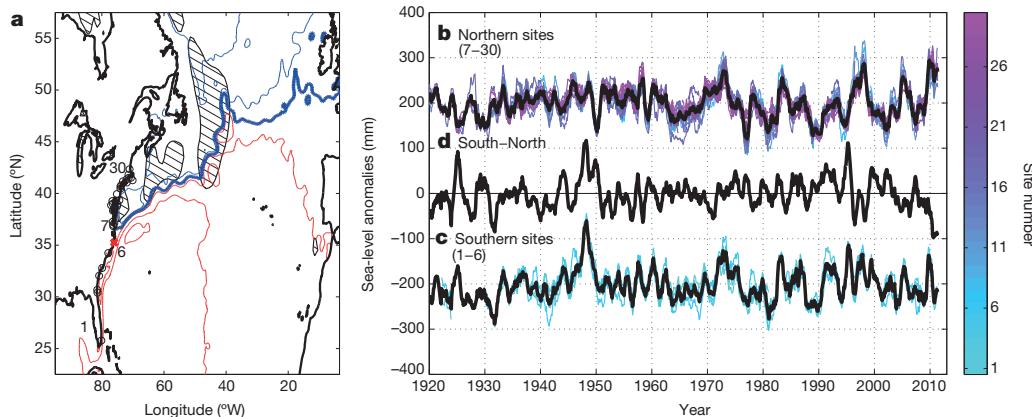


Figure 1 | Dynamic sea level and circulation along the western Atlantic seaboard. **a**, Negative (positive) mean dynamic topography contours in blue (red) indicate cyclonic (anticyclonic) geostrophic streamlines. Contour values in metres shown in Fig. 2. The zero contour (dark blue) marks the boundary between the subtropical and subpolar gyres. Hatched areas indicate warm sea surface temperature anomalies of greater than 0.5°C during the positive phase

signal is in the subpolar region (Fig. 1a), so we wish to show that, as a measure of ocean circulation, our sea-level index is related to heat transport into the subpolar gyre and consequently heat content changes there. Such a mechanism is supported by our model, in which the sea-level index leads the heat transport into the subpolar gyre at 40° N and, consequently, leads the heat content changes there (Extended Data Fig. 6).

Although we do not have observations of heat transport, we can relate our sea-level index directly to the heat content changes in the subpolar gyre since 1960. Figure 2a shows the accumulated sea-level index (blue curve), together with a direct estimate of the heat content in the area in the depth-weighted temperature anomaly in the top 500 m between 40° N and 60° N (black line). Heat content trends are similar throughout the upper 1,000 m of the Atlantic, below which they reverse due to the depth structure of the Atlantic overturning circulation. The cool subpolar upper ocean of the 1970s and 1980s and subsequent warming in the 1990s is captured by the accumulated sea-level index, observationally supporting the hypothesis that circulation changes and not only air-sea fluxes were involved in these changes²⁸. For the purposes of statistical analyses, the timeseries have had a 7-year low-pass, Tukey filter applied to them, which is referred to by the prefix ‘7-year’ from here on. The 7-year sea-level index leads the 7-year rate of

of the AMO from 1995–2004 relative to from 1961–2012. **b, c**, Dynamic sea-level anomalies north (**b**; sites 7–30, +200 mm offset) and south (**c**; sites 1–6, −200 mm offset) of Cape Hatteras, with averages in black. **d**, The difference in sea level, southern minus northern average, defines our sea-level index for ocean circulation.

heat content change by 2 years with a maximum correlation of 0.58 (significant at the 95% level). The reason that the accumulated sea-level index leads the large rise in heat content from 40° N to 60° N in the early 1990s can be interpreted by looking at maps of the heat content anomaly evolution. Heat content builds downstream of the intergyre region from the mid-1980s to the mid-1990s (Fig. 2b). This heat content anomaly is then observed downstream in the subpolar gyre in the late 1990s and early 2000s (Fig. 2c), indicating that the sea-level index could provide an early indication of subpolar heat content change.

The first mode of atmospheric variability over the North Atlantic, the North Atlantic Oscillation (NAO), forces both buoyancy and wind-driven ocean circulation⁷ and, we believe, is the major forcing of the circulation in the intergyre region. The 7-year NAO is significantly correlated with ($r = 0.71$ at the 98% level) and leads the 7-year sea-level difference by approximately 1 year over the period 1950 to 2012. On extending the time period to 1920–2012, the correlation drops slightly but is still significantly correlated ($r = 0.61$ at the 98% level, Extended Data Fig. 7). The correlation between the sea-level difference and the NAO is higher and more significant than the correlation of the NAO with either the southern or northern sea-level (Fig. 1b, c) composites ($r = -0.5$ at the 86% level for the southern composite; $r = -0.43$ at the 70% level for the northern); this supports

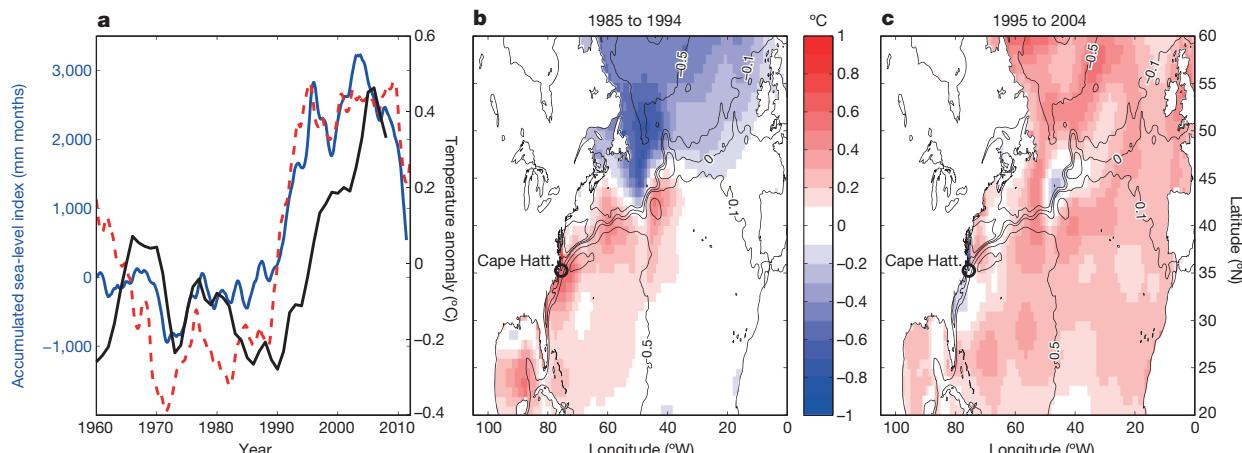


Figure 2 | Relating the sea-level circulation index to heat content changes. **a**, Accumulated sea-level index (nominally, in mm month) derived from accumulating the sea-level circulation index (blue), temperature anomaly in the upper 500 m of the subpolar North Atlantic from 40° to 60° N (black) and

accumulated NAO (red, dashed). **b**, Average temperature anomaly in the top 500 m for the periods 1985–94 relative to the average from 1958–2010. Contours of mean dynamic topography (metres) defined in Fig. 1a are overlaid for reference. **c**, Same as **b** but for the period 1995–2004.

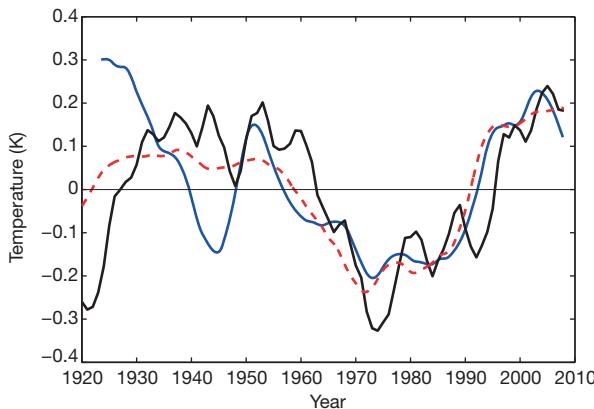


Figure 3 | Sea-level circulation index, the NAO and the AMO on multi-decadal timescales. Shown are the accumulated sea-level index (blue), which is representative of subpolar heat content evolution, the accumulated NAO (red, dashed) and the AMO (black). The heat content proxy and the accumulated NAO have been normalized. All timeseries have been 7-year low-pass filtered. The accumulated sea-level index and accumulated NAO have been detrended.

our hypothesis that the NAO forces the ocean circulation and consequently the ocean heat transport into the subpolar gyre.

In the past 90 years, the AMO has undergone three major transitions: warming in the mid-1990s and 1920s, and a cooling in the 1960s. From the early-1920s, when the tide gauge network along the east coast of North America was developed, robust comparisons of our sea-level index to the AMO are possible (Fig. 3). The accumulated sea-level index and the accumulated NAO are linearly detrended and capture much of the multi-decadal variation. The 7-year sea-level index leads the 7-year rate of change of the AMO by 2 years and is significantly correlated ($r = 0.51$, significant at the 96% level, Extended Data Fig. 8). This lead time of 2 years remains consistent when the timeseries is broken into 60-year blocks (Extended Data Table 2). In recent years, the sea-level index (Fig. 1d) indicates that the AMO is again transitioning to a negative phase, consistent with observations of a reduced overturning circulation⁸.

Using the sea-level difference between subtropical and subpolar gyres, we have developed and validated a proxy for ocean circulation in the intergyre region. This represents a mechanism for ocean heat transport to the subpolar gyre and heat content changes there. When observations exist, heat content changes have coincided with the major phase transitions of the AMO, confirming that ocean circulation plays a key role in decadal Atlantic variability. The ocean responds to NAO forcing with changes in ocean circulation: on decadal timescales, the ocean integrates NAO forcing and returns it to the atmosphere as the AMO. This is implicitly the Bjerknes compensation that had previously been seen in air-sea fluxes²⁹. The sea-level difference provides an indicator of ocean circulation changes that precede phase changes in the AMO, thus explaining why, as the positive AMO declines⁴, sea-level rise is accelerating north of Cape Hatteras^{9,10}. While Greenland ice sheet melt has been linked with accelerating sea-level rise in recent years, the fact that the period of accelerated sea-level rise from the 1950s to the 1970s¹⁰ (as well as the current period of sea-level rise) coincides with a declining AMO indicates that multi-decadal fluctuations in ocean circulation play a key role. In this framework, sea-level rise along the US east coast becomes entwined with the effects of the AMO on climate.

Online Content Methods, along with any additional Extended Data display items and Source Data, are available in the online version of the paper; references unique to these sections appear only in the online paper.

Received 3 July 2014; accepted 8 April 2015.

- Zhang, R. & Delworth, T. L. Impact of Atlantic multidecadal oscillations on India/Sahel rainfall and Atlantic hurricanes. *Geophys. Res. Lett.* **34**, L23708 (2006).

- Sutton, R. T. & Dong, B. Atlantic Ocean influence on a shift in European climate in the 1990s. *Nature Geosci.* **5**, 788–792 (2012).
- Goldenberg, S. B., Landsea, C. W., Mestas-Núñez, A. M. & Gray, W. M. The recent increase in Atlantic hurricane activity: causes and implications. *Science* **293**, 474–479 (2001).
- Chen, X. & Tung, K.-K. Varying planetary heat sink led to global-warming slowdown and acceleration. *Science* **345**, 897–903 (2014).
- Delworth, T. L. & Mann, M. E. Observed and simulated multidecadal variability in the Northern Hemisphere. *Clim. Dyn.* **16**, 661–676 (2000).
- Booth, B. B., Dunstone, N. J., Halloran, P. R., Andrews, T. & Bellouin, N. Aerosols implicated as a prime driver of twentieth-century North Atlantic climate variability. *Nature* **484**, 228–232 (2012).
- Marshall, J., Johnson, H. & Goodman, J. A study of the interaction of the North Atlantic Oscillation with ocean circulation. *J. Clim.* **14**, 1399–1421 (2001).
- Smeed, D. A. *et al.* Observed decline of the Atlantic Meridional Overturning Circulation 2004 to 2012. *Ocean Science* **10**, 39–38 (2014).
- Sallenger, A. H. Jr, Doran, K. S. & Howd, P. A. Hotspot of accelerated sea-level rise on the Atlantic coast of North America. *Nature Clim. Change* **2**, 884–888 (2012).
- Boon, J. D. Evidence of sea level acceleration at US and Canadian tide stations, Atlantic Coast, North America. *J. Coast. Res.* **28**, 1437–1445 (2012).
- Meinen, C. S., Baringer, M. O. & Garcia, R. F. Florida Current transport variability: an analysis of annual and longer-period signals. *Deep Sea Res. I* **57**, 835–846 (2010).
- Østerhus, S., Turrell, W. R., Jónsson, S. & Hansen, B. Measured volume, heat, and salt fluxes from the Atlantic to the Arctic Mediterranean. *Geophys. Res. Lett.* **32**, L07603 (2005).
- McCarthy, G. D. *et al.* Measuring the Atlantic Meridional Overturning Circulation at 26°N. *Prog. Oceanogr.* **130**, 91–111 (2015).
- Montgomery, R. Fluctuations in monthly sea level on eastern US coast as related to dynamics of western North Atlantic Ocean. *J. Mar. Res.* **1**, 165–185 (1938).
- Iselin, C. O. Preliminary report on long-period variations in the transport of the Gulf Stream system. *Pap. Phys. Oceanogr. Meteorol.* **3**, 1–40 (1940).
- Sturges, W. & Hong, B. Wind forcing of the Atlantic thermocline along 32°N at low frequencies. *J. Phys. Oceanogr.* **25**, 1706–1715 (1995).
- Wunsch, C. Mass and volume transport variability in an eddy-filled ocean. *Nature Geosci.* **1**, 165–168 (2008).
- Kanzow, T. *et al.* Basinwide integrated volume transports in an eddy-filled ocean. *J. Phys. Oceanogr.* **39**, 3091–3110 (2009).
- Bingham, R. J. & Hughes, C. W. Signature of the Atlantic meridional overturning circulation in sea level along the east coast of North America. *Geophys. Res. Lett.* **36**, L02603 (2009).
- Yin, J., Schlesinger, M. E. & Stouffer, R. J. Model projections of rapid sea-level rise on the northeast coast of the United States. *Nature Geosci.* **2**, 262–266 (2009).
- Johns, W. E. *et al.* Continuous, array-based estimates of Atlantic Ocean heat transport at 26.5°N. *J. Clim.* **24**, 2429–2449 (2011).
- Grist, J. P. *et al.* The roles of surface heat flux and ocean heat transport convergence in determining Atlantic Ocean temperature variability. *Ocean Dyn.* **60**, 771–790 (2010).
- Thompson, P. & Mitchum, G. Coherent sea level variability on the North Atlantic western boundary. *J. Geophys. Res.* **119**, 5676–5689 (2014).
- Ezer, T. Sea level rise, spatially uneven and temporally unsteady: why the US east coast, the global tide gauge record and the global altimeter data show different trends. *Geophys. Res. Lett.* **40**, 5439–5444 (2013).
- Hodson, D. L., Robson, J. I. & Sutton, R. T. An anatomy of the cooling of the North Atlantic Ocean in the 1960s and 1970s. *J. Clim.* **27**, 8229–8243 (2014).
- Häkkinen, S. & Rhines, P. B. Decline of subpolar North Atlantic circulation during the 1990s. *Science* **304**, 555–559 (2004).
- Bryden, H. L., King, B. A., McCarthy, G. D. & McDonagh, E. L. Impact of a 30% reduction in Atlantic meridional overturning during 2009–2010. *Ocean Sci.* **10**, 683–691 (2014).
- Robson, J., Sutton, R., Lohmann, K., Smith, D. & Palmer, M. D. Causes of the rapid warming of the North Atlantic Ocean in the mid-1990s. *J. Clim.* **25**, 4116–4134 (2012).
- Gulev, S., Latif, M., Keenlyside, N. S. & Koltermann, K. North Atlantic Ocean control on surface heat flux at multidecadal timescale. *Nature* **499**, 464–467 (2013).

Acknowledgements G.D.M. and D.A.S. are supported by the UK Natural Environment Research Council (NERC) RAPID-WATCH programme. I.D.H. was partly supported by the UK NERC consortium project iGlass (NE/I009906/1). J.P.G. and J.J.-M.H. are supported by NERC National Capability funding.

Author Contributions G.D.M. originated and developed the concept. I.D.H. provided the tide gauge data analysis. J.P.G. and J.J.-M.H. provided the numerical model analysis. D.A.S. carried out the statistical analysis. All authors contributed to the shaping and production of the manuscript.

Author Information Comma-separated data used in the manuscript are available to download from <http://bitly/1F7gtps>. Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to G.D.M. (gerard.mccarthy@noc.ac.uk).

METHODS

Data. Monthly mean sea-level records were obtained from the Permanent Service for Mean Sea-level (www.psmsl.org) for tide gauges stretching from Florida to Boston (locations 1 to 30, Extended Data Fig. 1). Linear trends were removed from each record. This removes the impact of glacial isostatic adjustment and other land subsidence effects, which have time periods of thousands of years and are known to affect tide gauges along this coastline. A 12-month low-pass filter removed the seasonal cycle. Southern (northern) composites of sea level were calculated by averaging records 1–6 (7–30). The meridional coherence of sea-level fluctuations is such that using just a single tide gauge results in an r.m.s. error of only 5 mm relative to the full composite. Finally, the sea-level index is simply the difference obtained by subtracting the northern from the southern sea-level composite. The high level of meridional coherence allows the interpretation of the sea-level gradient as this simple index.

Sources. Monthly NAO data from the National Center for Atmospheric Research “The Climate Data Guide: Hurrell North Atlantic Oscillation (NAO) Index (PC-based)” (<https://climatedataguide.ucar.edu/climate-data/hurrell-north-atlantic-oscillation-nao-index-pc-based>); monthly AMO index, based on the Kaplan sea surface temperature (SST) data set (from <http://www.esrl.noaa.gov/psd/data/timeseries/AMO/>); subsurface temperature data from the EN3 product (<http://www.metoffice.gov.uk/hadobs/en3/>); geostrophic velocity anomalies were produced and distributed by Aviso (<http://www.aviso.altimetry.fr/>) as part of the Ssalto ground processing segment. CNES-CLS09 Mean Dynamic Topography (v1.1 release) for the period 1993–99 was produced by the French Space Agency CNES.

Model validation. The multi-decadal oscillation of SSTs is most intense in the subpolar gyre (Fig. 1a). Modelling studies have shown that it is ocean heat transport into the subpolar gyre (here we choose 40° N) that controls the heat content of the subpolar upper ocean and consequently the SST. The concept here is that circulation in the intergyre region reflects the balance between warm subtropical water entering the subpolar gyre and colder subpolar water being recirculated within the gyre. We show that the sea-level gradient along the US east coast is a good proxy for this circulation (Extended Data Figs 4 and 5).

We can relate sea-level changes to ocean circulation in a reduced gravity geostrophic framework:

$$\mathbf{v} = \frac{\mathbf{g}'}{f} \mathbf{k} \times \nabla h$$

where \mathbf{v} is geostrophic velocity, \mathbf{k} is the unit vector in the vertical direction, h is sea level, \mathbf{g}' is reduced gravity and f is the Coriolis parameter. To estimate the transport in the intergyre region, previous studies have considered the sea-level difference between an onshore tide gauge and an offshore tide gauge, such as Bermuda. Reference 24, for example, relates the sea-level difference between Atlantic City and Bermuda to the Atlantic overturning circulation.

However, Bermuda is in the eddy-filled ocean interior¹⁷, which can disrupt spatially-coherent ocean transport signals. Our approach is to use sea-level estimates south of Cape Hatteras instead of an offshore sea-level estimate. Dynamic topography along the US east coast also decreases to the north across the intergyre boundary at Cape Hatteras much as it decreases from Bermuda to Atlantic City. However, measurements on the coast do not suffer the same contamination due to eddies as mid-ocean measurements¹⁸. Hence we estimate the transport along the intergyre boundary as:

$$v_{ig} \propto h_s - h_n$$

where the subscript ig refers to the intergyre region, s and n refer to south and north respectively. We can formulate the heat transport through a section straddling the intergyre boundary as:

$$HT_{ig} = \rho c_p \iint \Theta v_{ig} dA$$

where ρ is density, c_p is the specific heat capacity of seawater, Θ is conservative temperature and A is the area of the section considered. In this study we assume that the velocity fluctuations dominate the temperature fluctuations and so set the heat transport directly proportional to the intergyre velocity. This is an assumption that has proved true in direct heat transport estimates²¹. We note there is no dilemma in picking the location of the northern or southern points as the meridional coherence of sea-level fluctuations allows us to use a simple average of all sea-level records from Miami Beach to Cape Hatteras (Cape Hatteras to Boston) for h_s (h_n).

In terms of upper ocean heat content, the heat transported in this intergyre region has a profound impact on the subpolar gyre. This is because warm water may be transferred from the upper waters of the subtropics to the subpolar gyre

whereas subpolar water can only enter the subtropics at depth (traditionally in the deep western boundary current). Therefore we relate the heat transport into the subpolar gyre and heat content of the upper waters of the subpolar gyre to the transport in the intergyre region:

$$HT_{40^\circ N} \propto h_s - h_n$$

For exactly the reason that we need to use tide gauges as a proxy for heat transport, we cannot validate the conceptual model directly due to the lack of direct observations. However, a global eddy-permitting (1/4°) ocean model (ORCA-025) provides the framework to investigate these balances. The heat transport into the subpolar gyre has previously been shown in this model to be the dominant factor in setting upper ocean temperature in the subpolar gyre²². Here, we reproduce this result, showing that the accumulated heat transport across 40° N captures the major decadal fluctuations in heat content of the subpolar gyre (Extended Data Fig. 6). We can use these heat transport measurements to validate our circulation index. At this resolution there are shortcomings in the representation of the Gulf Stream path: the Gulf Stream overshoots at Cape Hatteras and separates from the US coast too far north. However, we take account of this in choosing the northern and southern sea-level points so that they straddle the separation point. Also, despite the model being eddy-permitting rather than eddy-resolving, it does generate mesoscale variability. This is seen when including an offshore sea-level measurement (such as Bermuda) in a sea-level circulation index. Such an index fails to reflect the large scale circulation. This effect would be expected to be even larger in an eddy-resolving model. Extended Data Fig. 5 shows that the model-derived sea-level index projects onto the intergyre velocities in a similar manner to the observed sea-level index. Extended Data Fig. 6 shows the accumulated sea-level difference compared with the accumulated heat transport across a section near 40° N and the volume averaged temperature of the upper 500 m of the subpolar gyre (40° N to 60° N). The sea-level difference is significantly correlated with the heat transport into the subpolar gyre ($r = 0.62$) and leads by 5 years (as in the main text, we report statistics on unaccumulated timeseries).

Code availability. Matlab code is available to download from <http://bit.ly/1F7gt�>.

Statistical analysis. Cross-correlations are calculated using annually averaged data after first removing the mean and linear trend from each variable. Two approaches are used to quantify the uncertainty in the correlation. First, we calculated the parameter

$$T = \sqrt{(N-2)} \frac{r}{\sqrt{1-r^2}}$$

where r is the correlation and N is the number of samples. The distribution of T is assumed to have a t -distribution with $N - 2$ degrees of freedom when the samples are not autocorrelated. This is used with a one-sided test to estimate the likelihood that the correlation has not occurred by chance (that is, the certainty with which we can reject the null hypothesis). Our data are autocorrelated and the number of independent samples (degrees of freedom) is therefore smaller than N . To calculate the effective number of degrees of freedom we follow ref. 30 by evaluating the autocorrelation of each variable and the estimate N as

$$N_{eff} = N_{obs} \frac{(1 - a_1 a_2)}{(1 + a_1 a_2)}$$

where N_{eff} is the degrees of freedom, N_{obs} is the number of observations and a_1 , a_2 are the values of the autocorrelations at a lag of one year. We evaluated N_{eff} over the longest time for each variable and then used the lowest value for all correlations. For the shorter time series N_{eff} was reduced in proportion to the length of the series. Degrees of freedom are reported in Extended Data Table 1.

In a second approach we applied the non-parametric method described in ref. 31. A large number (we used 10,000) of simulated time series are constructed from the Fourier transform of one of the original data series by preserving the modulus of each Fourier component but changing the phase to a random value between 0 and 2π . The distribution of correlations between these random series and the second variable was then calculated. The percentage of simulated correlations that are less than the observed correlation indicates the confidence that the true correlation is greater than zero. Because we are considering lagged correlations we modify the technique of ref. 31 so that for each simulated time series we evaluate the maximum of cross-correlation across all lags rather than the correlation at zero lag only. This provides a more stringent test of confidence.

To estimate the uncertainty in the time lag of the maximum correlation we used the times at which the correlation was equal to the maximum value less the

standard deviation of correlations derived from the simulated time series. The results are summarized in Extended Data Table 1.

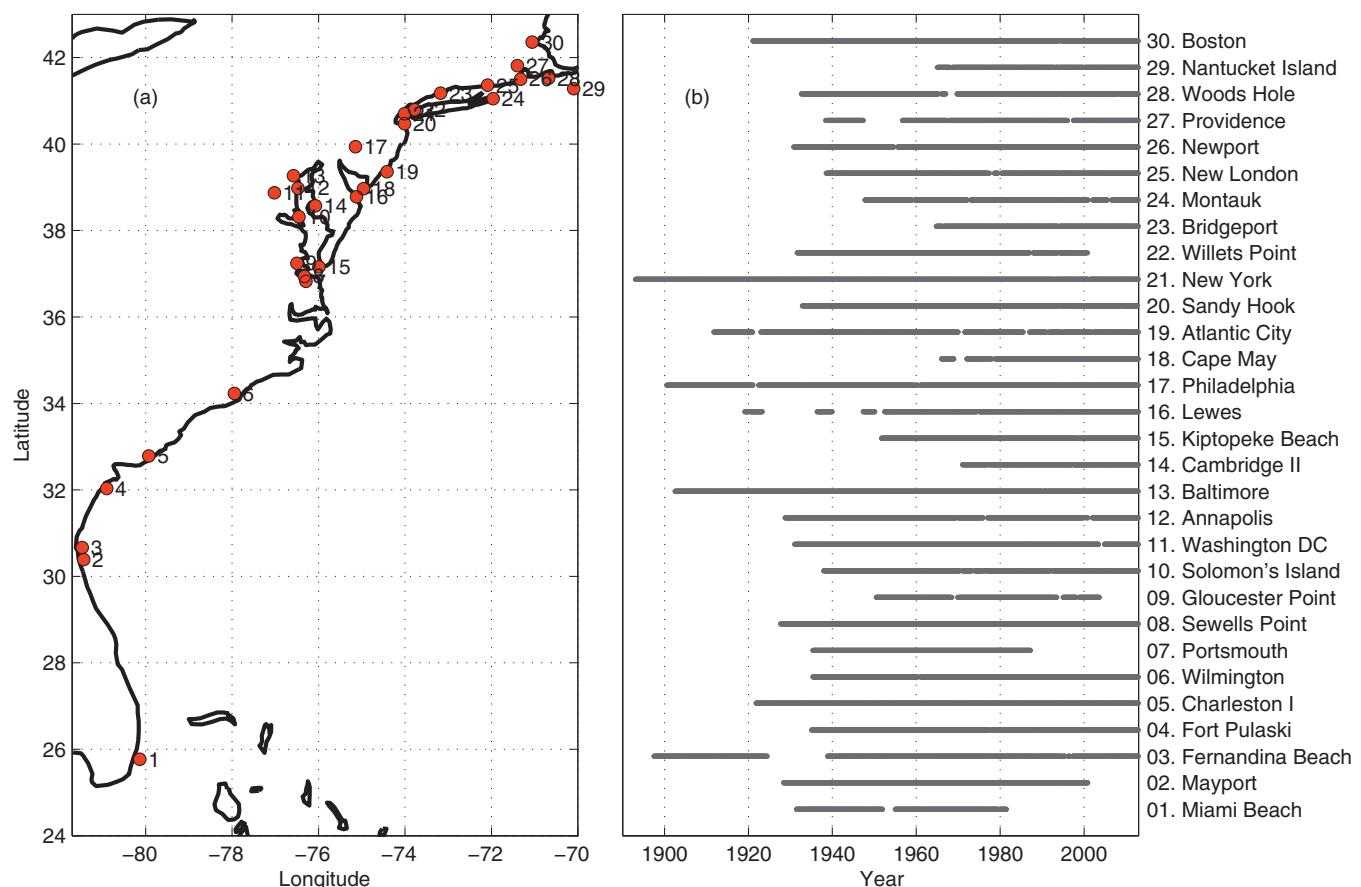
We have also evaluated the correlation over shorter periods to determine if the lag has remained constant over time. Results from three overlapping 60-year periods are shown in Extended Data Table 2. For each the correlation is a maximum when sea-level difference leads the differentiated AMO by 2 to 3 years.

The text refers to both accumulated and unaccumulated timeseries. Accumulation of zero mean timeseries constrains the beginning and end of the accumulated timeseries to zero. To avoid this arbitrary constraint, we report all our statistics on unaccumulated timeseries. As mentioned, for the purposes of

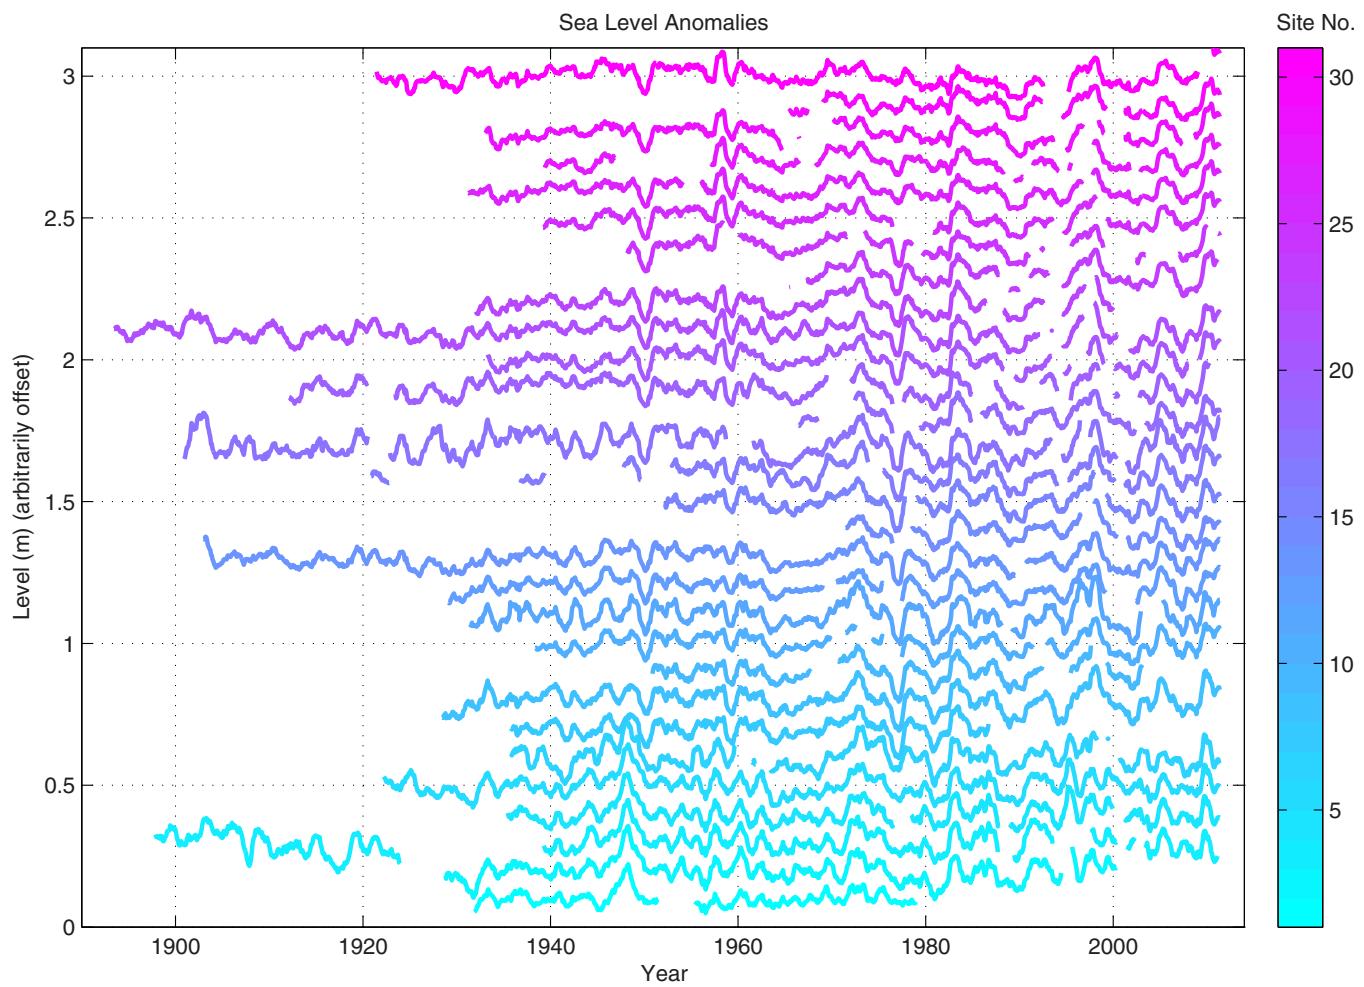
statistical analyses, the timeseries have had a 7-year, Tukey filter applied to them, which is referred to in the text with the prefix '7-year'.

Sample size. The timespan of the study was the maximum for which all of the necessary data were available. Therefore no statistical methods were used to pre-determine sample size as we used all the samples available to us.

30. Bretherton, C. S., Widmann, M., Dymnikov, V. P., Wallace, J. M. & Bladé, I. The effective number of spatial degrees of freedom of a time-varying field. *J. Clim.* **12**, 1990–2009 (1999).
31. Ebisuzaki, W. A method to estimate the statistical significance of a correlation when the data are serially correlated. *J. Clim.* **10**, 2147–2153 (1997).

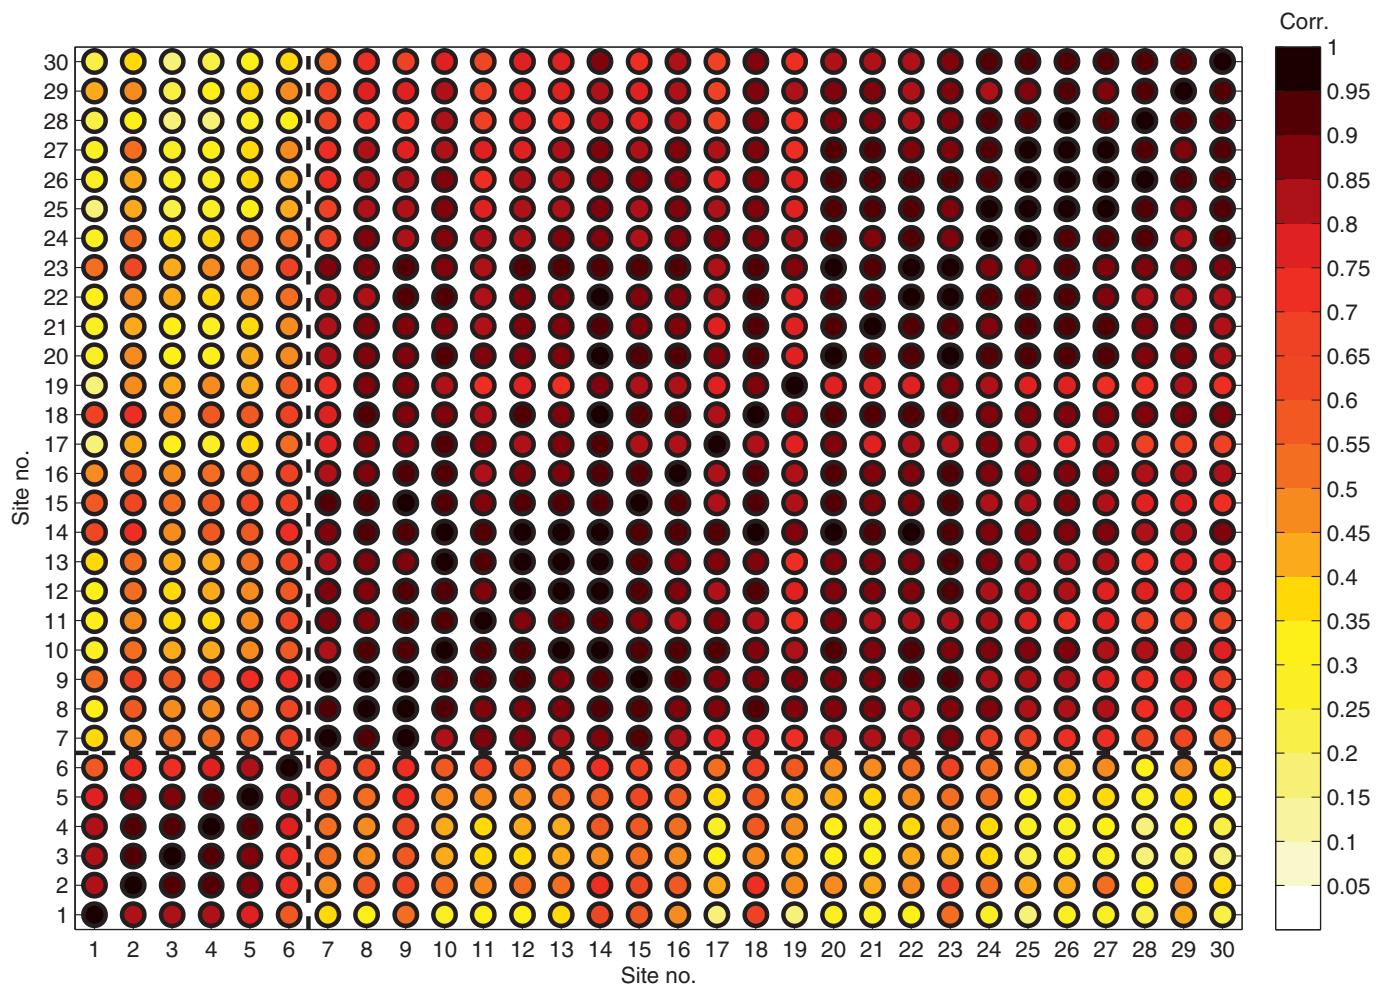


Extended Data Figure 1 | Tide gauges used in this study. **a**, Locations and **b**, temporal coverage of the tide gauges used in this study.

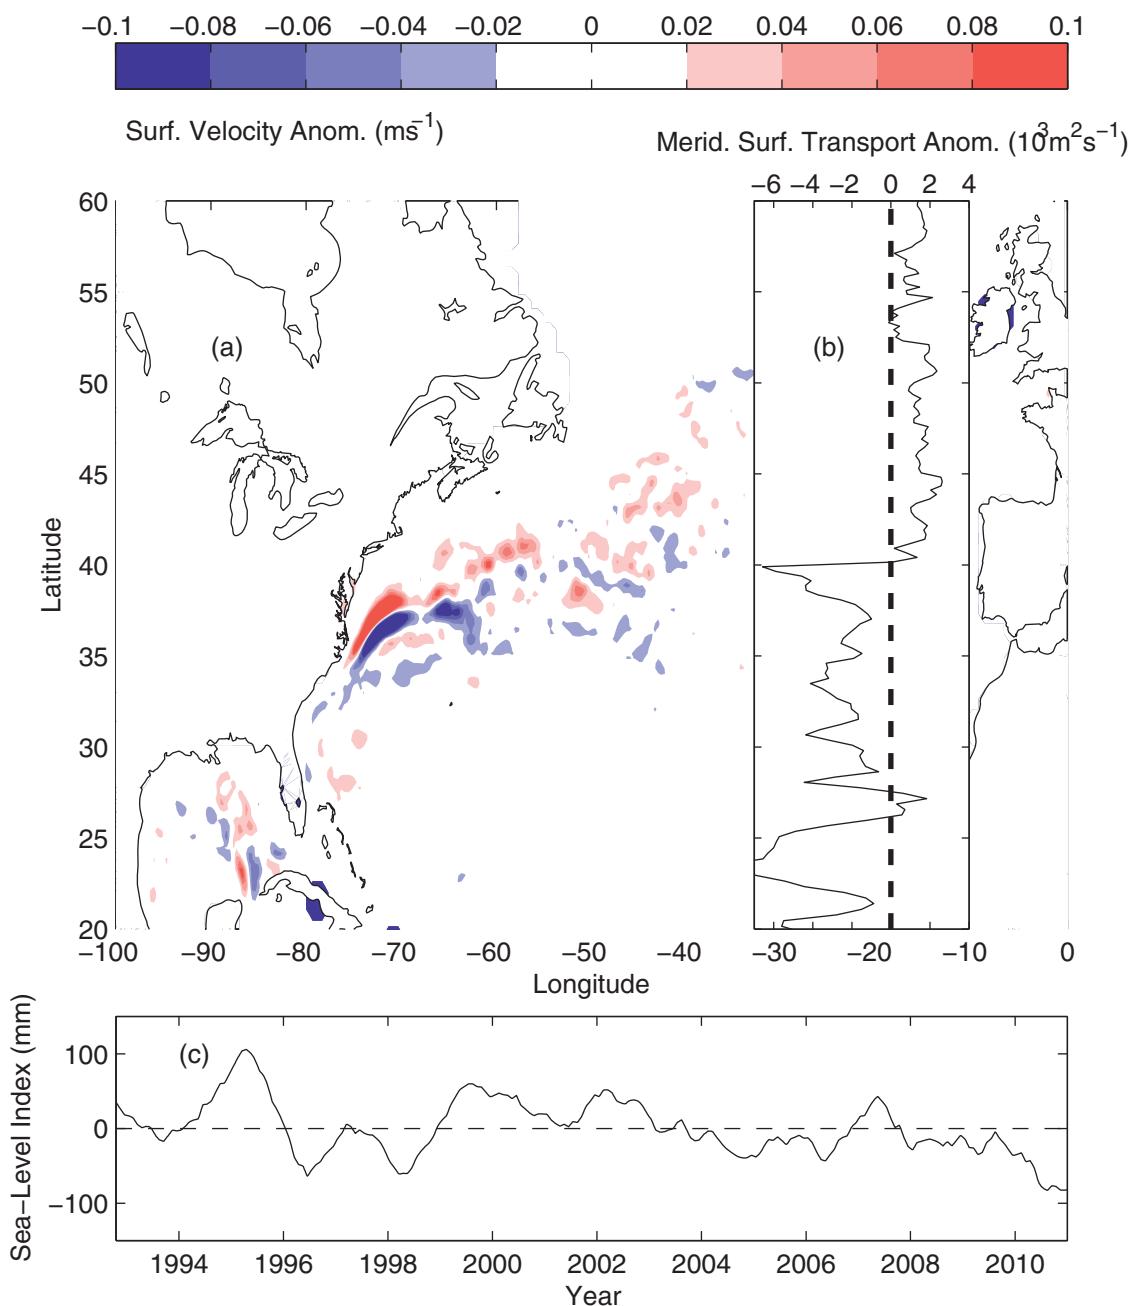


Extended Data Figure 2 | Dynamic sea-level anomalies from the 30 stations used in this study. Linear trends were removed from each record. This removes the impact of glacial isostatic adjustment and other land subsidence effects, which have time periods of thousands of years and are known to affect

tide gauges along this coastline. A seasonal cycle was removed using a 12-month boxcar filter. From 1920, there are multiple tide gauges both north and south of Cape Hatteras, so this is when we begin our study.

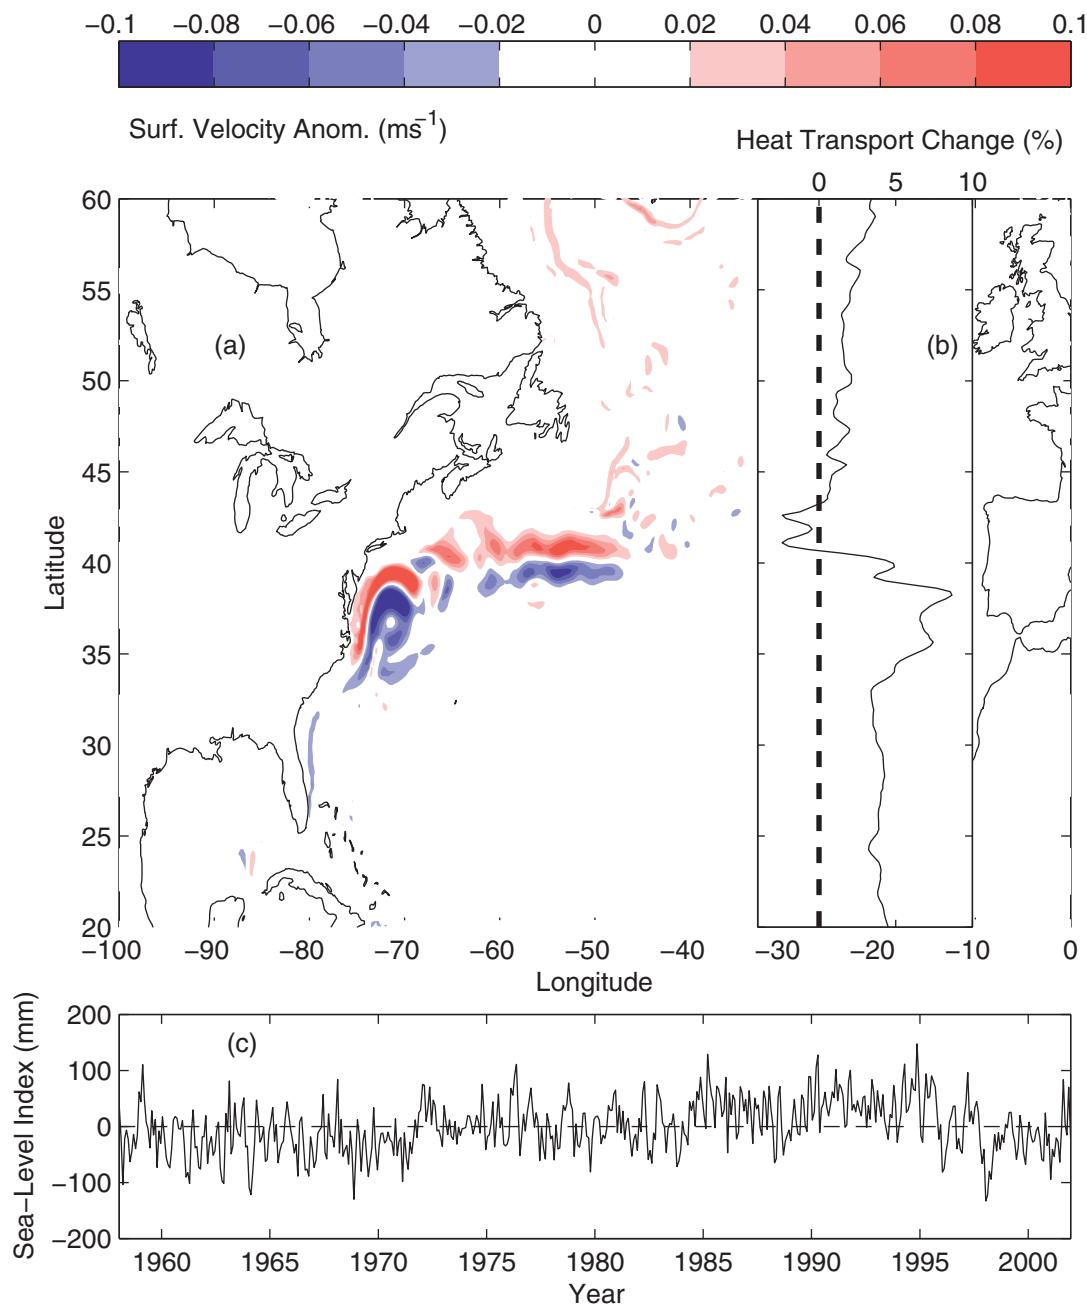


Extended Data Figure 3 | Correlation of tide gauges along the US east coast relative to one another. The dashed line indicates the location of Cape Hatteras. There is high correlation between tide gauges grouped north and south of Cape Hatteras.



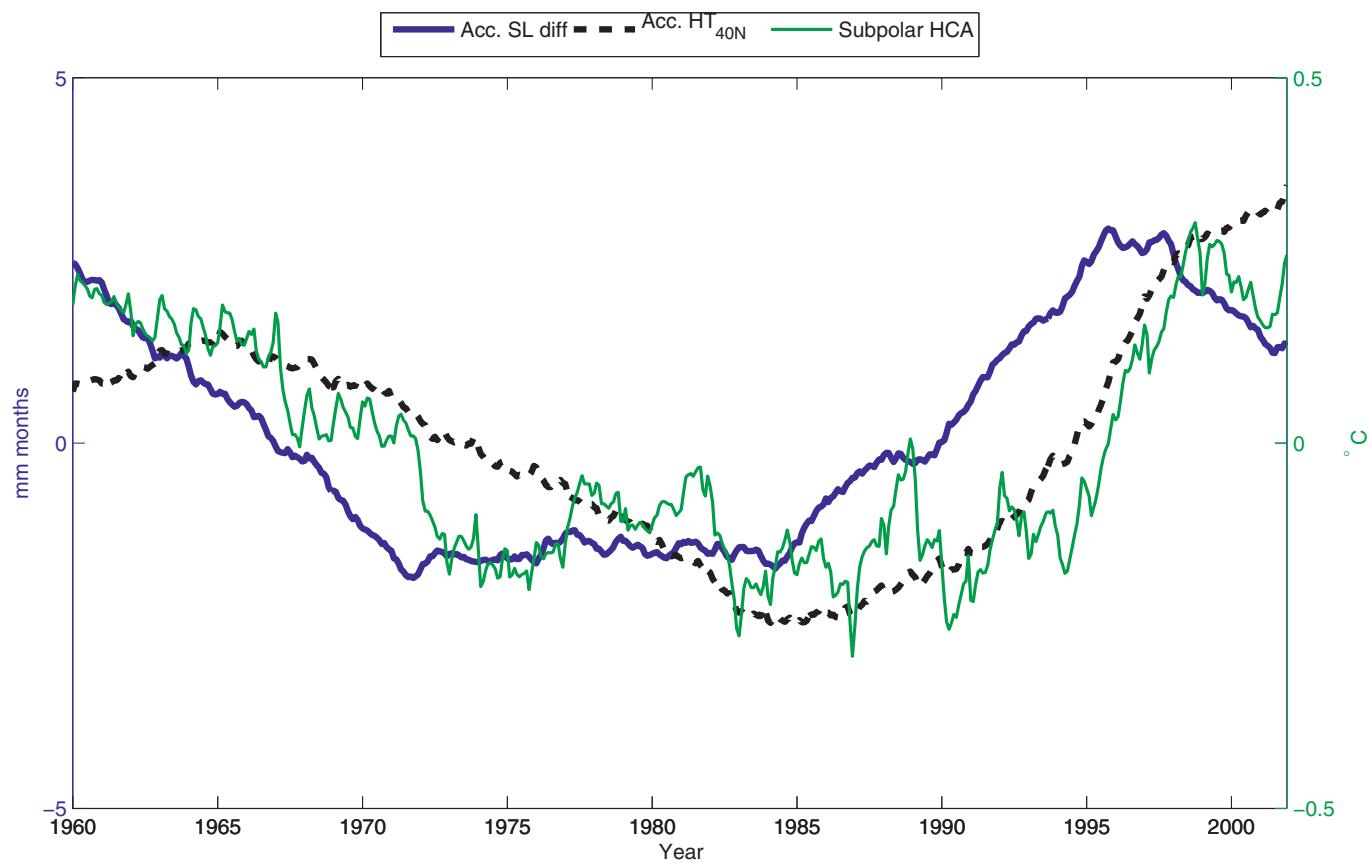
Extended Data Figure 4 | Surface velocity anomaly when the sea-level index is positive. **a**, Magnitude (m s^{-1}) and **b**, zonally integrated meridional velocity anomalies ($10^3 \text{ m}^2 \text{s}^{-1}$) for the time period 1993 to 2011, corresponding to when **(c)** the sea-level index is positive. A positive sea-level index is associated

with a more northerly circulation in the intergyre region and increased surface flow into the subpolar gyre. Velocities are geostrophic surface velocities derived from satellite altimetry.



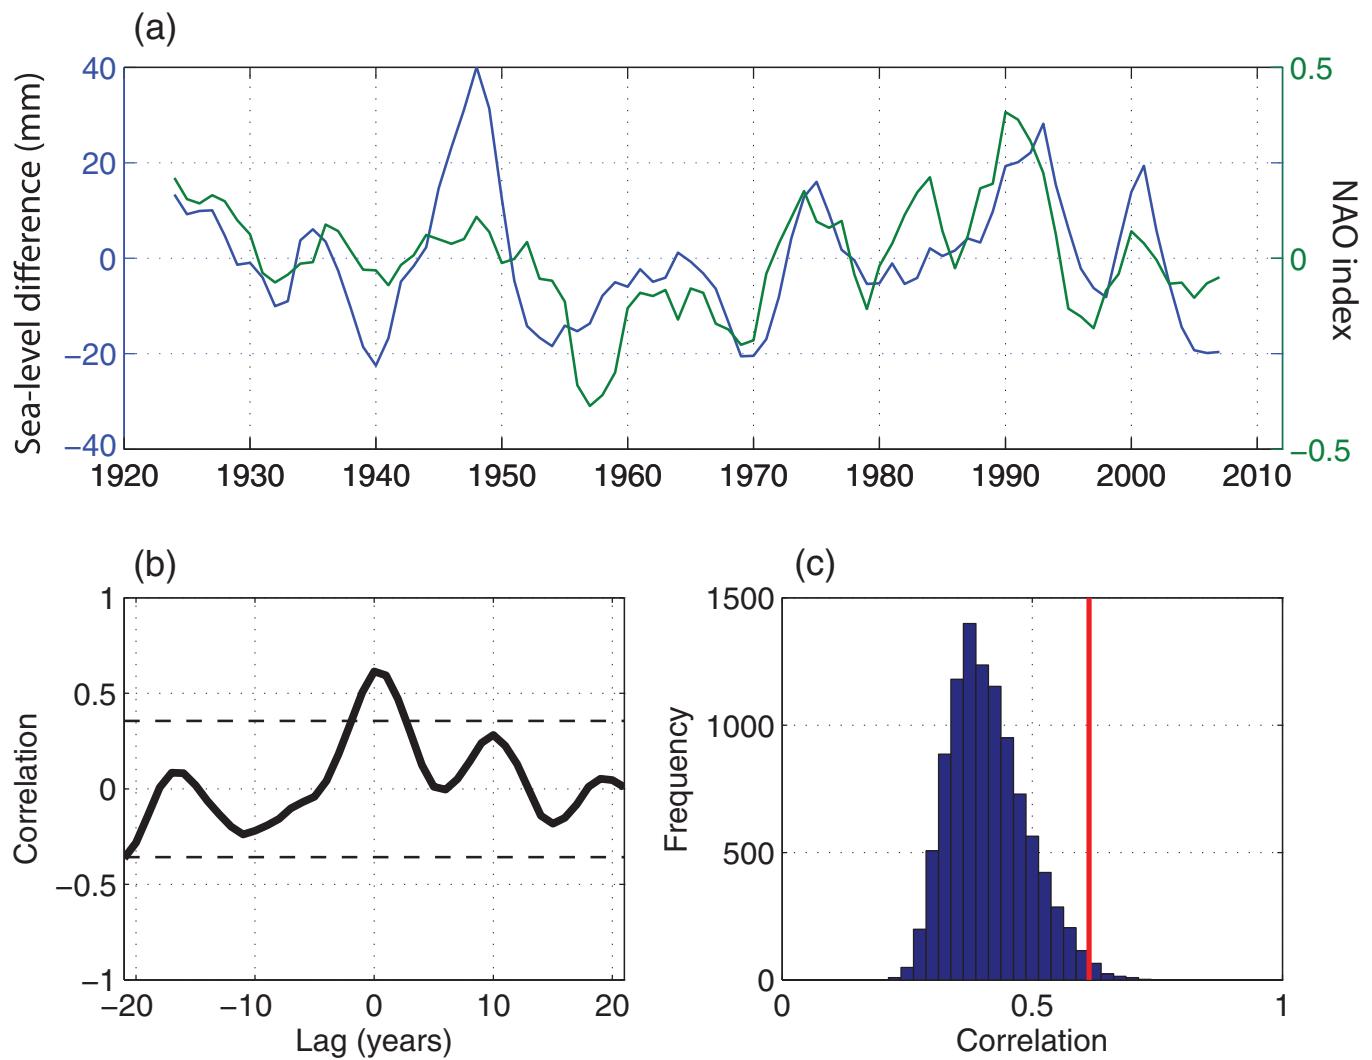
Extended Data Figure 5 | Model-derived surface velocity anomaly magnitude when the model-based sea-level index is positive. Similar to observed velocities, positive indices are associated with more northerly circulation in the intergyre region. **a**, Surface velocity magnitude (m s^{-1}) and **b**, percentage of meridional heat transport change (%) for the time period 1958

to 2001, corresponding to when (c) the model-derived sea-level index is positive. Similar to the satellite observations, a positive sea-level index is associated with a more northerly circulation in the intergyre region. Meridional heat transport change in both subtropical and subpolar gyres is positive when the sea-level index is positive.



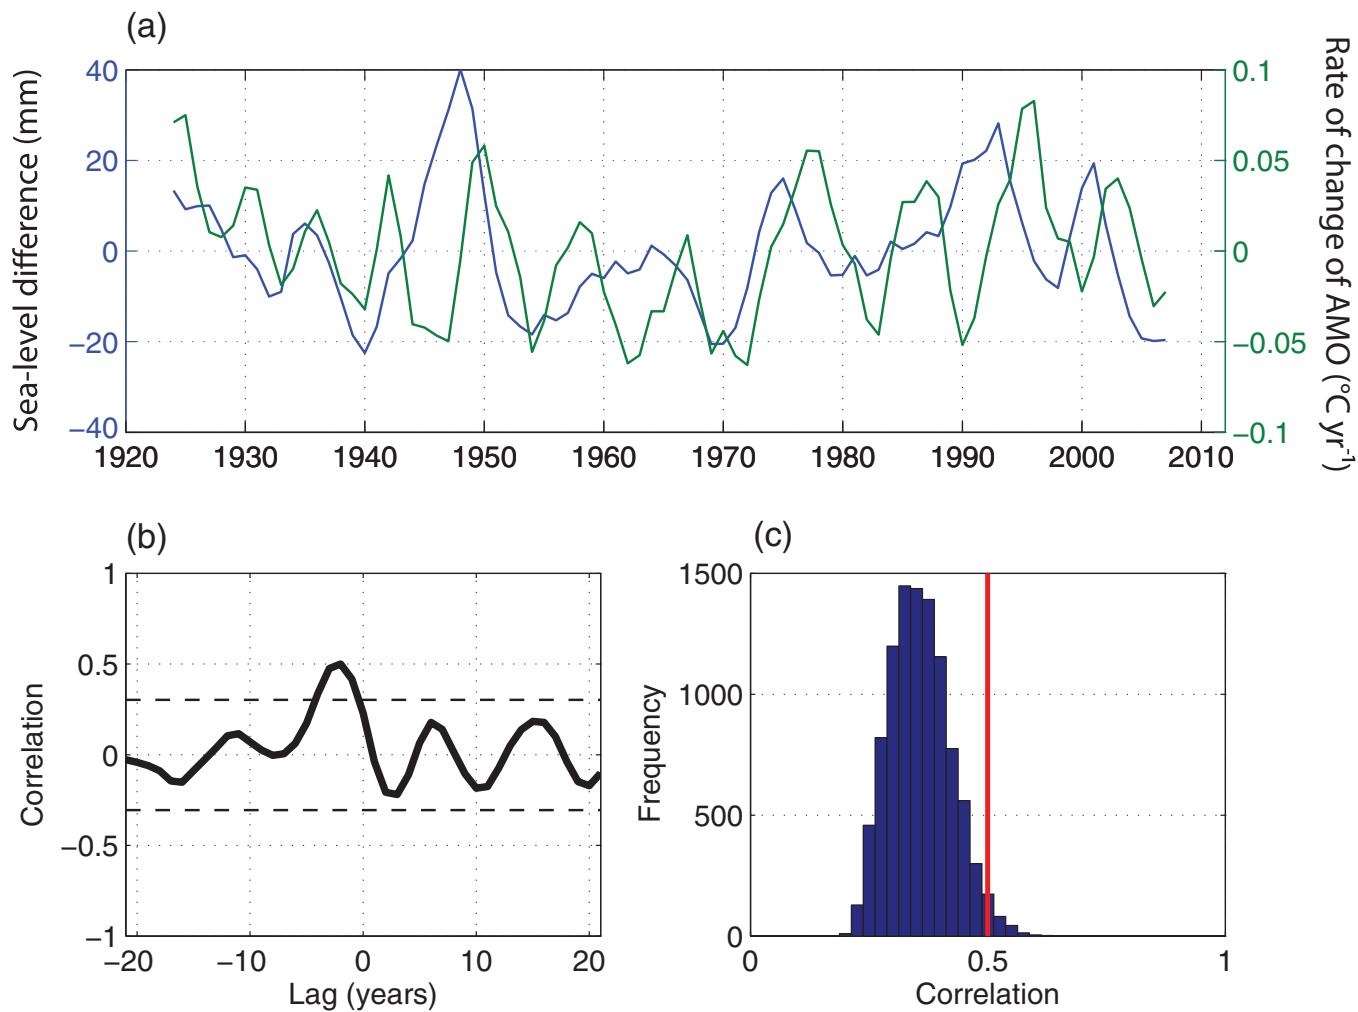
Extended Data Figure 6 | Model-derived sea-level index, heat transport and subpolar heat content. The accumulated sea-level index (Acc. SL diff, blue, in mm months) leads the accumulated heat transport into the subpolar

gyre across 40° N (Acc. HT_{40N} , black, normalized units). The heat transport into the subpolar gyre dominates the top 500 m temperature anomaly (Subpolar HCA, green, $^{\circ}\text{C}$) in the subpolar gyre.



Extended Data Figure 7 | Relationship between sea-level index and the NAO. **a**, 7-year sea-level difference (blue, cm) and 7-year NAO (green, normalized units). **b**, Lagged correlations between the two quantities. **c**, Scrambled correlation tests. The histogram indicates the typical correlations

that would be expected from randomly generated timeseries with similar spectral properties to the original timeseries. The red line indicates the maximum correlation between the two timeseries.



Extended Data Figure 8 | Relationship between sea-level index and the rate of change of the AMO. **a**, 7-year sea-level difference (blue) and rate of change of the AMO (green). **b**, Lagged correlations between the two quantities. **c**, Scrambled correlation tests. The histogram indicates the typical correlations

that would be expected from randomly generated timeseries with similar spectral properties to the original timeseries. The red line indicates the maximum correlation between the two timeseries.

Extended Data Table 1 | Sea level, NAO and rates of change of the AMO statistics

Var X	Var Y	Filt (yrs)	Time interval	DoF	Corr	Sig. % (t-stat)	Sig. % (Scrm)	RMS of rand corr	Lag at max corr (yrs)	Estimated range of lag (yrs)
B-A	Di HC	7	1950-2012	9	0.58	95	98	0.22	-2	-3 to 0
B-A	NAO	7	1950-2012	9	0.71	98	98	0.29	1	-1 to 2
NAO	Di HC	7	1950-2012	9	0.41	86	84	0.18	-2	-4 to -1
B-A	NAO	7	1920-2012	13	0.61	98	99	0.21	0	-1 to 2
B	NAO	7	1920-2012	13	-0.50	95	86	0.23	-11	-13 to 7
A	NAO	7	1920-2012	13	-0.43	91	70	0.22	1	-3 to 4
B-A	Di AMO	7	1920-2012	13	0.51	96	98	0.18	-2	-4 to -1
NAO	Di AMO	7	1920-2012	13	0.58	98	98	0.21	-4	-5 to -2

Correlation, lags and significance between the southern (B) and northern (A) sea-level indices. First variable (Var X) and second variable (Var Y) are indicated in the first two columns. In the second column, Di refers to the rate of change, and HC refers to subpolar heat content from 40° to 60° N. All timeseries are filtered with a 7-year low-pass filter (Filt). Degrees of Freedom (DoF), correlation (Corr) and significance (Sig.) are shown. Significance is determined by a t-statistic (t-stat) or a scrambled test (Scrm). The RMS of a random correlation (column 9) together with the lag at maximum correlation (column 10) allows determination of a range for the lag estimate (column 11).

Extended Data Table 2 | Sub-sampled statistics

From	To	Correlation	Lag (yrs)	Lag range (yrs)
1920	2012	0.5	-2	-4 to -1
1920	1980	0.36	-1	-3 to 0
1936	1995	0.46	-3	-4 to -1
1952	2012	0.54	-2	-4 to 0

Correlation, lags and lag range of sea-level index and the rate of change of the AMO over various time periods to investigate the consistency of the lags. Time periods under consideration are indicated in the first two columns.

Diverse coupling of neurons to populations in sensory cortex

Michael Okun^{1,2,3}, Nicholas A. Steinmetz^{1,2,3,4}, Lee Cossell^{2,5}, M. Florencia Iacaruso^{2,5}, Ho Ko²†, Péter Barthó⁶†, Tirin Moore⁴, Sonja B. Hofer^{2,5}, Thomas D. Mrsic-Flogel^{2,5}, Matteo Carandini³§ & Kenneth D. Harris^{1,2,6}§

A large population of neurons can, in principle, produce an astronomical number of distinct firing patterns. In cortex, however, these patterns lie in a space of lower dimension^{1–4}, as if individual neurons were “obedient members of a huge orchestra”⁵. Here we use recordings from the visual cortex of mouse (*Mus musculus*) and monkey (*Macaca mulatta*) to investigate the relationship between individual neurons and the population, and to establish the underlying circuit mechanisms. We show that neighbouring neurons can differ in their coupling to the overall firing of the population, ranging from strongly coupled ‘choristers’ to weakly coupled ‘soloists’. Population coupling is largely independent of sensory preferences, and it is a fixed cellular attribute, invariant to stimulus conditions. Neurons with high population coupling are more strongly affected by non-sensory behavioural variables such as motor intention. Population coupling reflects a causal relationship, predicting the response of a neuron to optogenetically driven increases in local activity. Moreover, population coupling indicates synaptic connectivity; the population coupling of a neuron, measured *in vivo*, predicted subsequent *in vitro* estimates of the number of synapses received from its neighbours. Finally, population coupling provides a compact summary of population activity: knowledge of the population couplings of n neurons predicts a substantial portion of their n^2 pairwise correlations. Population coupling therefore represents a novel, simple measure that characterizes the relationship of each neuron to a larger population, explaining seemingly complex network firing patterns in terms of basic circuit variables.

The cortex represents its computations through the joint activity of multiple neurons. This activity can be remarkably diverse even among neighbouring neurons, belonging to the same morphological and laminar cell class. In sensory cortex, neighbouring neurons not only respond to diverse stimulus features^{6–8}, but also use diverse strategies to encode information. For example, mean firing rate differs by orders of magnitude across neurons^{9–11}, and it appears to constitute an invariant property of each cell, persisting across multiple stimulus conditions and spontaneous activity¹². We asked whether there are other invariant dimensions that characterize the diversity of firing of cortical neurons. Ideally, such dimensions would not only help explain the complex patterns of activity produced by cortical populations, but also relate directly to underlying circuit variables.

To characterize how different cortical neurons relate to large-scale firing patterns, we considered their relationship to the ‘population rate’, that is, the summed activity of all neurons in the recorded population at any moment in time^{13–16}. We recorded first from populations of 20–80 neurons in deep layers of area V1 in awake head-fixed mice using multi-site silicon probes (Fig. 1a). Consistent with previous reports^{17–19}, population rate was not strictly controlled by sensory stimulation, but showed

coherent fluctuations even spontaneously in the absence of sensory stimuli (Fig. 1b, c), closely tracking the simultaneously recorded local field potential (LFP, Fig. 1d).

During spontaneous activity, neighbouring neurons differed widely in their coupling to the population. The spike-triggered population rate typically had a single peak centred close to zero (Fig. 1e and Extended Data Fig. 1), whose height was large for some neurons and small or even reversed for others. The variation in population coupling across neurons was continuous, and neurons did not appear to fall into discrete classes of distinct coupling. However, this diversity in coupling represented a robust property of individual cells; computing coupling from separate data segments yielded highly consistent results (Extended Data Fig. 2a). A similar result was seen when correlating spikes with the LFP (Fig. 1f and Extended Data Fig. 2b). These variations in population coupling were not trivially explained by differences in mean firing rate; shuffling the spike times across neurons and time, while preserving each neuron’s mean rate and the population rate distribution¹⁴, destroyed the diversity in population coupling between neurons (Fig. 1g and Extended Data Fig. 2c). Moreover, diversity in population coupling was not restricted to mouse V1; it was also observed in monkey area V4 (Fig. 1h) and in rat (*Rattus norvegicus*) auditory cortex (Extended Data Fig. 3).

Population coupling differed both between and within cell classes. On the basis of spike shape²⁰, we classified ~13% of the cells in mouse V1 as narrow-spiking putative interneurons (Extended Data Fig. 4a), with the remainder primarily pyramidal cells. Narrow-spiking cells tended to have higher population coupling than wide-spiking cells (medians of 0.84 and 0.50; $P < 10^{-6}$). Nevertheless, this difference between classes was substantially smaller than the variability within each class (Extended Data Fig. 4). Strongly coupled neurons were more likely to exhibit bursting and to have a lower firing rate. However, even cells with similar firing rate and bursting dynamics could show great variation in coupling (Extended Data Fig. 4b–d).

We next asked to what degree population coupling can explain the patterns of spontaneous correlated activity in a population. To this end, we generated random population activity patterns, constrained so that each neuron’s population coupling and mean rate matched the original data, as did the distribution of population rate over time (Fig. 2a). We found that pairwise correlations in the resulting synthetic activity resembled those observed in the recorded activity (Fig. 2b). The accuracy of this prediction was higher when the population rate showed stronger variance, indicating that the single statistic of coupling was able to capture the correlations induced by globally coordinated fluctuations (Fig. 2c). If coupling was not included, however, the performance of the model was substantially impaired (Fig. 2c, d).

These results indicate that much of the pairwise correlation in the population is explained by the coupling of each neuron to population

¹UCL Institute of Neurology, University College London, London WC1N 3BG, UK. ²Department of Neuroscience, Physiology and Pharmacology, University College London, London WC1E 6DE, UK. ³UCL Institute of Ophthalmology, University College London, London EC1V 9EL, UK. ⁴Howard Hughes Medical Institute and Department of Neurobiology, Stanford University, Stanford, California 94305-5125, USA. ⁵Biozentrum, University of Basel, Klingelbergstrasse 50/70, CH-4056 Basel, Switzerland. ⁶Center for Molecular and Behavioral Neuroscience, Rutgers University, 197 University Avenue, Newark, New Jersey 07102, USA. †Present addresses: Lui Che Woo Institute for Innovative Medicine and Chow Yuk Ho Technology Center for Innovative Medicine, Faculty of Medicine, the Chinese University of Hong Kong, Shatin, New Territories, Hong Kong, China (H.K.); MTA-TTK NAP-B Sleep Oscillations Research Group, Research Center for Natural Sciences, Hungarian Academy of Sciences, Magyar tudósok körútja 2, 1117 Budapest, Hungary (P.B.).

§These authors jointly supervised this work.

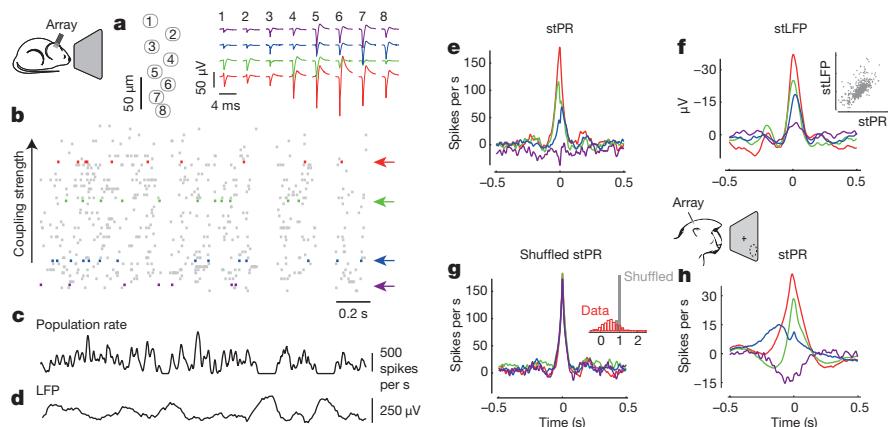


Figure 1 | Neighbouring neurons differ markedly in population coupling during spontaneous activity. **a**, Schematic of a single shank of silicon electrode array, and spike waveforms of four example wide-spiking neighbour neurons measured with the array in deep layers of V1 of an awake mouse. **b**, Population raster of spontaneous activity in 66 neurons recorded from the whole array. Cells are arranged vertically in order of population coupling. Arrows indicate the four example neurons shown in **a**. **c**, Population rate measured by summing all the spikes detected on the entire array. **d**, LFP measured on a shank adjacent to that on which the example neurons were recorded (LFP waveforms were similar across shanks). **e**, Spike-triggered population rate (stPR) for the four example neurons. The spike train of each

neuron was excluded from the population rate before computing its stPR. **f**, The spike-triggered local field potential (stLFP) for the four example cells (inverted for ease of comparison) resembles their stPR (shown in **e**). Inset, normalized magnitudes of stPR and stLFP (see Methods) are highly correlated across cells ($\rho = -0.71, P < 10^{-100}$, rank correlation, $n = 431$ neurons).

g, Differences in population coupling disappear after shuffling spikes in a manner that preserves each neuron's mean firing rate and the population rate. Inset, population couplings in the actual spike trains (red) and after shuffling (grey), for neurons from all experiments. **h**, stPR of four example neurons simultaneously recorded in primate area V4, computed as in **e**.

rate, with coupling strengths that vary between neurons. The underlying model is parsimonious, requiring only order n parameters to predict order n^2 pairwise correlations. Moreover, the model is intuitive, involving procedures among the simplest in neuroscience—summing the activity of multiple neurons, and correlating the spike train of each neuron with the result. To assess whether more advanced procedures would yield different results, we used a variant of latent variable analysis designed for discrete spike count data^{4,21} to obtain the weights of

individual neurons to the first detected factor. Reassuringly, these weights were highly correlated with population coupling (Extended Data Fig. 5a); latent variable analysis found the same basic structure as our simple coupling model.

The ability of the model to predict correlations may appear surprising given that it operates without knowledge of the neurons' sensory tuning. In primary sensory areas, neurons with similar sensory selectivity show stronger stimulus-independent correlations^{1,3,5,22}, and we observed a

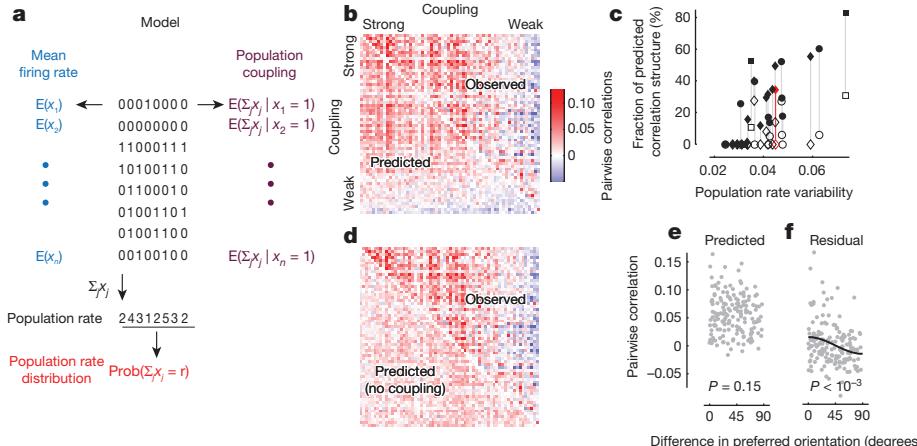


Figure 2 | A simple model based on population coupling predicts the structure of pairwise correlations in a cortical population. **a**, The model generates random spike patterns subject to three constraints: that the population coupling of each neuron, the mean firing rate of each neuron, and the distribution of the population rate must match those in the original data. **b**, Random activity generated by the model produces pairwise correlations that are similar to those measured in the original spike trains ($n = 67$ units in one experiment; correlations computed in 20-ms bins). The upper triangle shows observed pairwise correlations, and the lower one shows pairwise correlations predicted by the model. Neurons are arranged in order of population coupling. The values on the diagonal (all 1s) have been removed. Similarity of observed and predicted correlations is indicated by the symmetry of the upper and lower triangles. **c**, Percentage of explainable correlation structure predicted, as a function of the variability of population rate (filled symbols, see Methods). The model captures pairwise correlations, but only in

experiments in which the population rate fluctuates. It cannot predict them when population rate is mostly constant (a highly desynchronized cortical state). Recordings were obtained from mouse V1 in wakefulness (diamonds) or under anaesthesia (circles), or from A1 of awake rat (squares), all spontaneous activity; note that a variety of states is observed in all conditions. Open symbols show predictions of a model that ignores population coupling. The example experiment in **b** is shown in red. **d**, Same as **b** for predictions made without using population coupling. Such predictions fail to capture the structure of pairwise correlations (open markers in **c**). **e**, The model cannot predict a relationship between similarity of preferred orientation and spontaneous pairwise correlations ($P = 0.15$, Pearson correlation). **f**, As a result, this correlation is retained in the residual pairwise correlations obtained by subtracting the modelled from actual correlations ($\rho = 0.26, P < 10^{-3}$, Pearson correlation), indicating that the predictions of coupling and orientation sum linearly. The black line in **f** shows regression on $\cos(2\Delta\theta)$.

significant relationship between orientation preference similarity and spontaneous correlations (Pearson correlation $\rho = 0.26, P < 10^{-3}$, for $n = 163$ pairs of cells with orientation selectivity index >0.6). The coupling model could not predict this relationship; indeed, its predictions of correlations were independent of orientation preference similarity ($P = 0.15$, Fig. 2e). However, population coupling represented a much better predictor of pairwise correlations than orientation preference similarity ($\rho = 0.63, P < 10^{-9}$); moreover, the two predictors are independent, representing complementary predictions that sum to explain the fine structure of pairwise correlations within a population (Fig. 2f).

Population coupling was invariant to visual stimulation, and showed little relation to basic visual properties. We measured population coupling across neurons during responses to natural movies, and found that it closely resembled that obtained during spontaneous activity (Fig. 3a, b and Extended Data Fig. 6a, b). Differences in coupling, therefore, reflect invariant properties of each neuron, rather than differences in the responses of each neuron to a specific set of stimuli. Both strongly and weakly coupled cells could respond to natural movies, showing reliable modulation by visual stimuli (Fig. 3c, d). In responses to gratings, moreover, population coupling showed no relationship to orientation selectivity, spatial frequency preference, or linearity of spatial summation (Extended Data Fig. 6c–e).

However, population coupling predicted the overall increase in mean firing rate caused by a visual stimulus. Strongly coupled cells exhibited a larger increase in mean firing rate during visual stimulation, both in responses to natural stimuli (Fig. 3c–e; $\rho = 0.38, P < 10^{-15}, n = 431$) and in responses to drifting gratings ($\rho = 0.32, P = 2 \times 10^{-6}, n = 217$; Extended Data Fig. 6f). Thus population coupling—a quantity defined on a time scale of tens of milliseconds, and measurable during spontaneous activity—predicts the temporally extended, nonspecific elevation in firing rate seen during sensory stimulation.

These correlates of population coupling do not prove that the relationship between population rate and individual cells' firing is causal. To test causality, we used transgenic mice that express channelrhodopsin-2 in a subset of layer 5 pyramidal neurons²³. Light stimulation (0.5-s pulses at ~ 2 mW per mm^2 using a blue LED) strongly increased network activity (Fig. 3f). A control experiment (brief laser pulses at 100 mW per mm^2) indicated that in most cells this increase was driven synaptically. Only a few cells ($\sim 5\%$) reliably responded at short enough latencies (<5 ms) to indicate direct optogenetic depolarization. While strongly coupled cells substantially increased their firing in response to LED stimulation, activity in weakly coupled cells showed little increase, or even decreased (Fig. 3g). Therefore, the diverse coupling of individual neurons to population rate reflects differences in the causal influence of the population on these cells.

These observations suggest that more strongly coupled neurons might receive stronger synaptic input from neighbouring neurons. Consistent with this view, intracellular *in vivo* measurements indicated that population coupling measured from a neuron's spike train resembles the coupling measured from the neuron's subthreshold membrane potential, which mostly reflects synaptic inputs (Extended Data Fig. 7; see also ref. 17).

To investigate whether population coupling is related to synaptic connectivity, we studied cortical populations that were first imaged *in vivo* and then recorded *in vitro*^{24,25}. Population activity in superficial cortical layers was measured during presentation of natural movies and images in lightly anaesthetized mice. As in deeper layers, these cells showed a wide heterogeneity in population coupling (Fig. 4a, b). Coupling was not related to a cell's spatial location (Fig. 4b and Extended Data Fig. 8a), and we observed no correlation between *in vivo* coupling and intrinsic properties such as resting potential, input resistance and spike threshold (Extended Data Fig. 8b–d). Population coupling was again consistent between the halves of each recording, even when these halves contained different sets of stimuli ($\rho = 0.85, P < 10^{-9}, n = 4,215$ neurons from 15 experiments). Next, we analysed synaptic connectivity

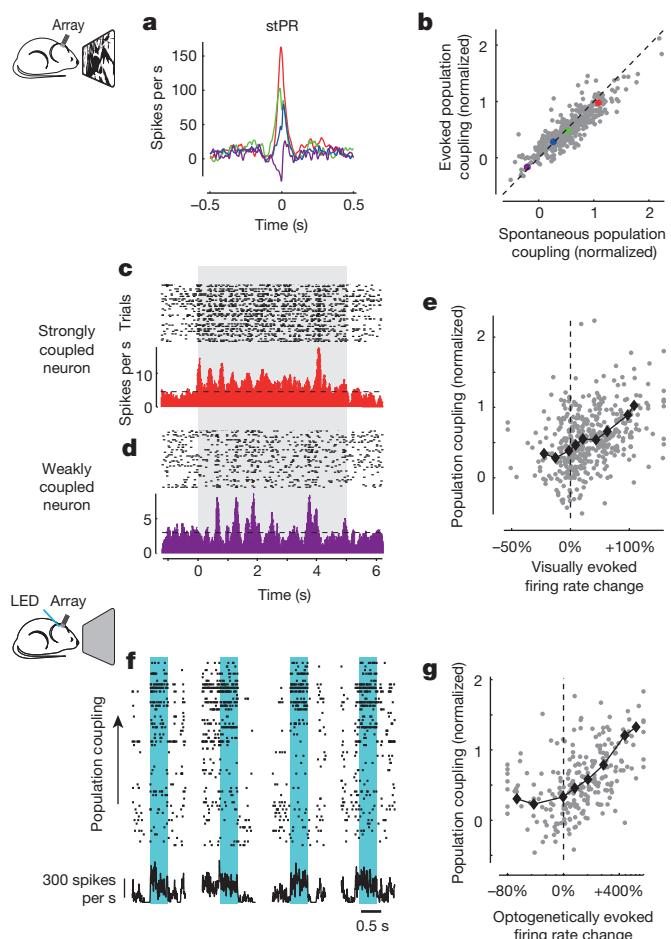


Figure 3 | Population coupling under natural and optogenetic stimulation conditions. **a**, Spike-triggered population rate (stPR) for the four example neurons in Fig. 1a–g during responses to a natural movie. The curves are similar to those measured during spontaneous activity (Fig. 1e). **b**, Comparison of population coupling during spontaneous and evoked activity, across cells and experiments ($\rho = 0.88, P < 10^{-100}$, rank correlation). **c**, Spike rasters (one row per presentation of a natural movie) and corresponding firing rate for a strongly coupled neuron (red neuron in a). Dashed line indicates baseline firing rate. Shaded area shows the duration of the movie clip. **d**, Same as c for a weakly coupled neuron (purple neuron in a). **e**, The increase in mean firing rate of a cell in response to natural movie presentations (relative to baseline) correlates with population coupling measured during spontaneous activity ($\rho = 0.38, P < 10^{-15}$, rank correlation, $n = 431$ neurons from 13 recordings in 8 animals). Black diamonds, running median. Points outside the x-axis range appear at the border for display purposes. **f**, Population rasters showing activity of deep-layer V1 neurons during four example trials (out of 75 in total), where the network was optogenetically driven by blue light in a mouse expressing channelrhodopsin-2 sparsely in layer 5. Neurons are sorted by their population coupling during spontaneous activity. Shaded area shows the duration of optogenetic stimulation. **g**, Change in mean firing rate evoked by optogenetic stimulation correlates with population coupling measured during spontaneous activity ($\rho = 0.51, P < 10^{-100}$, rank correlation, $n = 237$ neurons). Points outside the x-axis range appear at the border for display purposes. Black diamonds, running median.

using paired *in vitro* whole-cell recordings of the same neurons (Fig. 4c). We focused on pyramidal cells, identified by somatodendritic morphology when filled with Alexa dye, and by a regular-spiking firing pattern.

These paired recordings revealed a significant correlation between a neuron's population coupling and its probability of receiving synaptic input from its neighbours (Fig. 4d). According to a logistic regression fit, a pyramidal cell whose population coupling was one standard deviation below the average had a probability of 0.14 to receive a connection from a neighbouring neuron, whereas for a pyramidal neuron whose

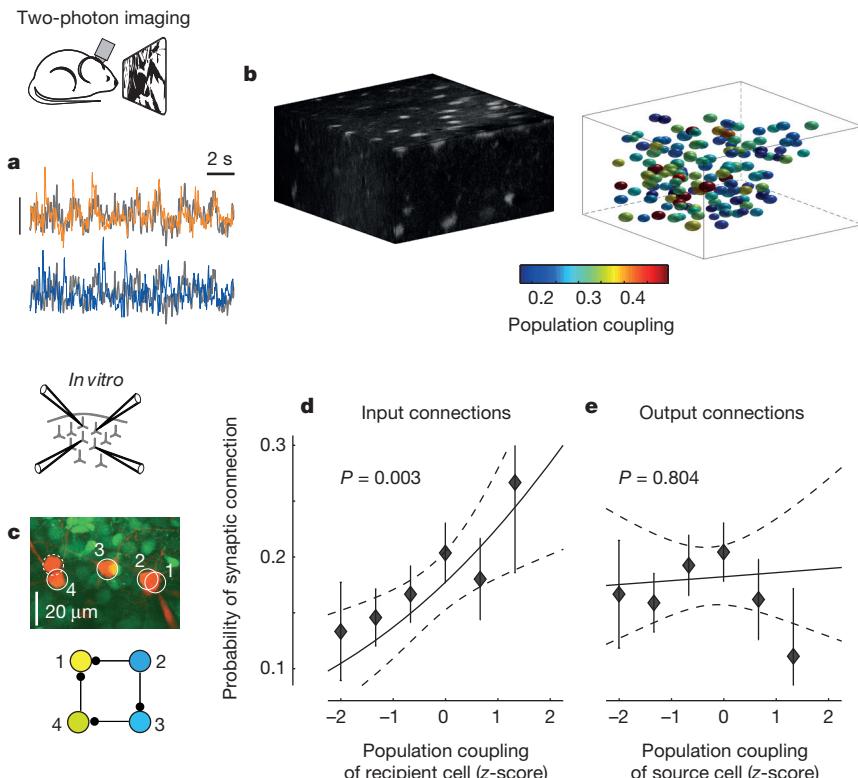


Figure 4 | Neurons with strong population coupling receive more synaptic inputs from their neighbours. **a**, Neurons in superficial layers of mouse V1 were bulk-loaded with Oregon Green BAPTA-1 dye and their activity recorded using two-photon imaging during presentation of natural movies and images. Population coupling was assessed as the correlation between each cell's calcium signal with the summed signal of all other neurons. The two coloured traces show a segment of activity from a strongly and a weakly coupled neuron (orange and blue, respectively), each superimposed on the averaged population activity (grey). Scale bar, 20% $\Delta F/F$ for each single neuron, 5% for population average. **b**, Left, structural scan of an imaged volume measuring $\sim 260 \times 260 \times 56 \mu\text{m}$. Right, pseudocolour representation of population coupling for each of 147 neurons in the volume. **c**, Synaptic connectivity of

a subset of the imaged neurons was later assessed using simultaneous *in vitro* whole-cell recordings. Top, four example pyramidal cells (solid white circles), recorded *in vitro* together with an additional fast-spiking interneuron (dashed circle) that was excluded from later analysis. Bottom, four synaptic connections were found between these four pyramidal neurons, shown here coloured by their *in vivo* population coupling. The two weakly coupled neurons (blue) received zero or one input, while the strongly coupled neuron (yellow) received two inputs. **d**, Logistic regression estimate of probability to receive a synaptic connection as a function of the population coupling of the target (dashed black lines, 95% confidence intervals; error bars, mean \pm standard error for binned data). **e**, As in **d** for outgoing connections.

coupling was a standard deviation above the mean this probability increased to 0.23 (a 65% increase). In contrast, there was no correlation between a neuron's population coupling and the probability that it provided synaptic outputs to neighbouring cells (Fig. 4e). These results are consistent with a correlation between population coupling and input connection probability as high as $\rho = 0.65$ (Extended Data Fig. 9).

These data suggest that diversity in mean input connectivity may form the circuit mechanism for diversity in population coupling. The diversity in mean input connectivity that we have observed is separate from the tendency of cortical pyramidal cells to receive more synaptic connections from cells with similar sensory tuning^{24,25}. Thus, if the network is stimulated non-specifically (for example by optogenetic activation of a random neuronal subset, Fig. 3f,g), the effects of this non-specific stimulation, amplified by recurrent cortical connections, will be largest for neurons with strong mean input connectivity (Extended Data Fig. 10a, b). To make these intuitive ideas more precise, we extended an established modelling framework²⁶, and mathematically analysed a model cortical circuit in which different excitatory neurons have different mean input probability (see Supplementary Information). This model reproduced many of our experimental findings, including the relationship between population coupling and mean input connectivity (Extended Data Fig. 10c and Fig. 4d), the structure of spontaneous pairwise correlations (Extended Data Fig. 10d and Fig. 2b), and the additivity of correlations predicted by sensory tuning similarity and by population coupling strength (Extended Data Fig. 10e-h and Fig. 2e, f).

If neurons with stronger population coupling are more likely to be driven by non-specific excitation, they may show greater modulation by non-sensory factors, which are believed to be conveyed to sensory cortex by diffuse, non-specific inputs^{15,19}. To test this hypothesis, we analysed population recordings made from area V4 of monkeys performing a cued-saccade task²⁷. Consistent with the hypothesis, strongly coupled neurons increased their firing rate during saccade preparation into their receptive field (Fig. 5a). Weakly coupled neurons often showed the opposite effect (Fig. 5b). Overall, there was a marked correlation between a neuron's population coupling and its modulation by saccade preparation ($\rho = 0.37$, $P = 10^{-9}$, Fig. 5c).

In conclusion, population coupling constitutes a previously unappreciated dimension characterizing the relationship of individual neurons to population activity. Strongly coupled neurons (choristers) are more strongly activated during multiple conditions that non-specifically increase the activity of their local network: not only sensory stimuli, but also spontaneous fluctuations, polysynaptic optogenetic stimulation, and top-down modulation. Conversely, weakly coupled neurons (soloists) are more immune to these population-wide events. Population coupling differs both across and within cell classes; it remains to be determined whether this within-class diversity reflects further subdivisions of these classes (such as pyramidal cells with different long-range axonal targets^{28,29}), or continuous within-class variation in cellular parameters such as input connection probability. Moreover, a neuron's population coupling need not be fixed for life; neurons that rats learn to

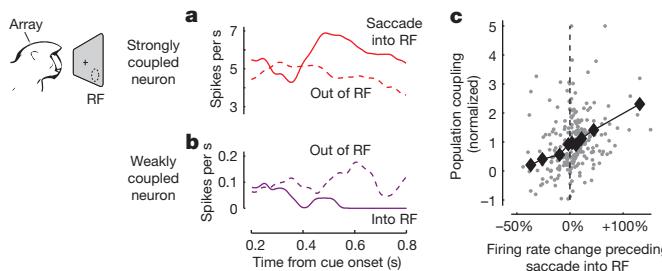


Figure 5 | Population coupling under top-down stimulation conditions.

a, Mean firing rate of an example neuron in primate V4 with strong population coupling (red neuron in Fig. 1h). This neuron showed higher firing rate (solid curve) while saccades were prepared into its receptive field (RF) than outside of it (dashed curve). b, An example neuron with weak population coupling (purple neuron in Fig. 1h) showed suppressed firing during saccade preparation into its receptive field. c, The change in firing rate of V4 neurons during saccade preparation into their receptive field (relative to saccade out of their receptive field) correlates with population coupling ($\rho = 0.37$, $P = 10^{-9}$, rank correlation, $n = 262$ neurons). Black diamonds, running median.

use while controlling a brain-machine interface show increased correlation with LFP³⁰, indicating increased population coupling, and thus suggesting an increase in mean synaptic input strength. This single, simple variable relating each neuron's participation in the population code to underlying circuit connectivity may prove critical to understanding cortical computation and plasticity.

Online Content Methods, along with any additional Extended Data display items and Source Data, are available in the online version of the paper; references unique to these sections appear only in the online paper.

Received 1 September 2014; accepted 30 January 2015.

Published online 6 April 2015.

1. Tsodyks, M., Kenet, T., Grinvald, A. & Arieli, A. Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* **286**, 1943–1946 (1999).
2. Yu, B. M. et al. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *J. Neurophysiol.* **102**, 614–635 (2009).
3. Luczak, A., Bartho, P. & Harris, K. D. Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron* **62**, 413–425 (2009).
4. Pfau, D., Pnevmatikakis, E. A. & Paninski, L. in *Advances in Neural Information Processing Systems* Vol. 26, 2391–2399 (Curran Associates, 2013).
5. Kenet, T., Arieli, A., Tsodyks, M. & Grinvald, A. in *23 Problems in Systems Neuroscience* (eds van Hemmen, J. L. & Sejnowski, T. J.) Ch. 9 (Oxford University Press, 2006).
6. Ohki, K., Chung, S., Ch'ng, Y. H., Kara, P. & Reid, R. C. Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. *Nature* **433**, 597–603 (2005).
7. Rothschild, G., Nelken, I. & Mizrahi, A. Functional organization and population dynamics in the mouse primary auditory cortex. *Nature Neurosci.* **13**, 353–360 (2010).
8. Martin, K. A. & Schroder, S. Functional heterogeneity in neighboring neurons of cat primary visual cortex in response to both artificial and natural stimuli. *J. Neurosci.* **33**, 7325–7344 (2013).
9. Hromádka, T., Deweese, M. R. & Zador, A. M. Sparse representation of sounds in the unanesthetized auditory cortex. *PLoS Biol.* **6**, e16 (2008).
10. Sakata, S. & Harris, K. D. Laminar structure of spontaneous and sensory-evoked population activity in auditory cortex. *Neuron* **64**, 404–418 (2009).

11. O'Connor, D. H., Peron, S. P., Huber, D. & Svoboda, K. Neural activity in barrel cortex underlying vibrissa-based object localization in mice. *Neuron* **67**, 1048–1061 (2010).
12. Buzsáki, G. & Mizuseki, K. The log-dynamic brain: how skewed distributions affect network operations. *Nature Rev. Neurosci.* **15**, 264–278 (2014).
13. Renart, A. et al. The asynchronous state in cortical circuits. *Science* **327**, 587–590 (2010).
14. Okun, M. et al. Population rate dynamics and multineuron firing patterns in sensory cortex. *J. Neurosci.* **32**, 17108–17119 (2012).
15. Harris, K. D. & Thiele, A. Cortical state and attention. *Nature Rev. Neurosci.* **12**, 509–523 (2011).
16. Tkačík, G. et al. Searching for collective behavior in a large network of sensory neurons. *PLOS Comput. Biol.* **10**, e1003408 (2014).
17. Okun, M., Naim, A. & Lampl, I. The subthreshold relation between cortical local field potential and neuronal firing unveiled by intracellular recordings in awake rats. *J. Neurosci.* **30**, 4440–4448 (2010).
18. Luczak, A., Bartho, P. & Harris, K. D. Gating of sensory input by spontaneous cortical activity. *J. Neurosci.* **33**, 1684–1695 (2013).
19. Zagha, E., Casale, A. E., Sachdev, R. N., McGinley, M. J. & McCormick, D. A. Motor cortex feedback influences sensory processing by modulating network state. *Neuron* **79**, 567–578 (2013).
20. Barthó, P. et al. Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. *J. Neurophysiol.* **92**, 600–608 (2004).
21. Macke, J. et al. in *Advances in Neural Information Processing Systems* Vol. 24, 1350–1358 (Curran Associates, 2011).
22. Kohn, A. & Smith, M. A. Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J. Neurosci.* **25**, 3661–3673 (2005).
23. Arenkiel, B. R. et al. *In vivo* light-induced activation of neural circuitry in transgenic mice expressing channelrhodopsin-2. *Neuron* **54**, 205–218 (2007).
24. Ko, H. et al. Functional specificity of local synaptic connections in neocortical networks. *Nature* **473**, 87–91 (2011).
25. Ko, H. et al. The emergence of functional microcircuits in visual cortex. *Nature* **496**, 96–100 (2013).
26. Douglas, R. J., Koch, C., Mahowald, M., Martin, K. A. C. & Suarez, H. H. Recurrent excitation in neocortical circuits. *Science* **269**, 981–985 (1995).
27. Steinmetz, N. A. & Moore, T. Eye movement preparation modulates neuronal responses in area v4 when dissociated from attentional demands. *Neuron* **83**, 496–506 (2014).
28. Yamashita, T. et al. Membrane potential dynamics of neocortical projection neurons driving target-specific signals. *Neuron* **80**, 1477–1490 (2013).
29. Chen, J. L., Carta, S., Soldado-Magraner, J., Schneider, B. L. & Helmchen, F. Behaviour-dependent recruitment of long-range projection neurons in somatosensory cortex. *Nature* **499**, 336–340 (2013).
30. Gulati, T., Ramanathan, D. S., Wong, C. C. & Ganguly, K. Reactivation of emergent task-related ensembles during slow-wave sleep after neuroprosthetic learning. *Nature Neurosci.* **17**, 1107–1113 (2014).

Supplementary Information is available in the online version of the paper.

Acknowledgements We thank L. Buesing for advice on latent variable analysis, C. Reddy and T. Sato for technical assistance, and M. Häusser for advice on the manuscript. This work was supported by the Wellcome Trust (S.B.H., T.D.M.-F., M.C., K.D.H.), Engineering and Physical Sciences Research Council (K.D.H.), the European Research Council (T.D.M.-F.), the Medical Research Council (L.C.), National Institutes of Health (EY014924, N.A.S. and T.M.) and the Simons Foundation (M.C. and K.D.H.). M.C. holds the GlaxoSmithKline/Fight for Sight Chair in Visual Neuroscience.

Author Contributions M.O. conceived the study and performed the *in vivo* electrophysiology experiments in mouse V1. M.O., N.A.S. and L.C. performed the analyses. N.A.S. and T.M. performed the experiments in primate V4. L.C., M.F.I., H.K., S.B.H. and T.D.M.-F. performed the imaging and *in vitro* experiments and contributed to data analyses. P.B. performed the experiments in rat A1. K.D.H. constructed the mathematical model. M.O., M.C. and K.D.H. designed the study and wrote the paper.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to M.O. (m.okun@ucl.ac.uk) or K.D.H. (kenneth.harris@ucl.ac.uk).

METHODS

Electrophysiological recordings. Recordings in mouse primary visual cortex were performed under licenses from the UK Home Office in accordance with the Animal (Scientific Procedures) Act 1986. In the electrophysiological experiments, recordings were made from male and female mice older than 6–7 weeks, of C57BL/6J or Thy1-ChR2 line 18 strains²³. Randomization was not required as all animals were treated similarly.

For recordings from awake mice, animals were first implanted with a custom-built head plate and recording chamber under isoflurane anaesthesia. After 3 days of recovery accompanied by Rimadyl treatment, the mice underwent two or three head-restraint acclimatization sessions. On the recording day, the animals were briefly anaesthetized with isoflurane and a craniectomy of $\sim 1.5 \times 1.5$ mm was made above the left primary visual cortex. The dura was resected with a 30G needle and the brain was covered with Ringer solution and Kwik-Cast (WPI Inc.). At least 1.5 h of recovery between this procedure and the beginning of the recording were allowed, a period that was more than sufficient for the animals to fully resume their normal behaviour in their home cage. In some animals, an additional recording was performed the following day; in this case the brain was again protected with Ringer solution and Kwik-Cast in between the two recording sessions. Thirteen recordings were made in eight animals, of which five recordings were in three Thy1-ChR2 line 18 mice, and the rest in C57BL/6J mice. Recordings were made with Buzsaki32 or A4x8 silicon probes (NeuroNexus technologies, Ann Arbor, Michigan), lowered to a depth of 550–1050 μm (median 700 μm) by a PatchStar manipulator (Scientifica, Uckfield, UK). Signals were amplified and stored for offline analysis using the Cerebus data acquisition system (Blackrock Microsystems, Salt Lake City, Utah). For electrophysiological recordings in anaesthetized animals (Fig. 2 and Extended Data Fig. 7), anaesthesia was induced with a mixture of fentanyl citrate, fluanisone and midazolam (0.8 mg kg^{-1} , 25 mg kg^{-1} and 12.5 mg kg^{-1} , respectively). A total of nine recordings from anaesthetized animals were analysed. In a few of these experiments, whole-cell recordings were performed in parallel with the silicon probe population recording. The whole-cell intracellular recordings were performed following the standard techniques for blind patching^{17,31,32}, with the patch pipette positioned ~ 200 –400 μm away from the plane of the silicon probe shanks.

For the electrophysiological data, the experimental setup was similar to the one described in ref. 33, except that mice were standing in a custom-built tube, instead of on a floating Styrofoam ball. In brief, visual stimuli were presented on two of the three available LCD monitors, positioned ~ 25 cm from the animal and covering a field of view of $\sim 120^\circ \times 60^\circ$, extending in front and to the right of the animal. Visual stimuli analysed in the present study consist of multiple presentations of natural movie video clips (taken from *The Life of Mammals*, BBC, London, UK) and drifting gratings (100% and 25% contrast, 2 Hz, 12 directions, spatial frequencies of 0.02, 0.04, 0.08 and 0.25 cycles per degree). For recordings of spontaneous activity, the monitors showed a uniform grey background.

Recordings in auditory cortex (Extended Data Fig. 3) were made in awake, head-fixed male rats, in accordance with protocols approved by the Rutgers University Animal Care and Use Committee. Experimental details were previously described in refs 3, 18, 34. In the recordings analysed here, silicon probes with four shanks, and four contacts in a tetrode configuration on each shank were used (NeuroNexus technologies, Ann Arbor, Michigan).

Recordings in primate area V4 were made in two male monkeys (*Macaca mulatta*) using 16-channel U-Probes (Plexon Inc., Dallas, Texas), in accordance with NIH Guide for the Care and Use of Laboratory Animals, the Society for Neuroscience Guidelines and Policies, and Stanford University Animal Care and Use Committee; full experimental details were previously described in ref. 27. In brief, as part of the behavioural task, a cue presented near the central fixation point indicated the direction of a saccade that would be necessary at the end of the trial to obtain a reward in the event of a change in the orientation of a peripheral stimulus.

Optogenetic stimulation. For *in vivo* electrophysiological recordings with optogenetic stimulation, we used Thy1-ChR2 line 18 mice, a transgenic line in which ChR2 is expressed in a sparse subset of cortical L5 pyramidal cells²³. In these mice, in addition to recording spontaneous and visually evoked activity, we also recorded activity in response to optogenetic stimulation consisting of 500 ms light pulses of $\sim 2 \text{ mW per mm}^2$ intensity, delivered by a 470 nm (blue) LED (M470F1, Thor Labs, Newton, New Jersey), and focused on a circular area of $\sim 1 \text{ mm}^2$ centred on the silicon probe insertion site.

To estimate the proportion of recorded neurons directly activated by ChR2, we performed a control experiment, in which 20- and 40-Hz trains of five 2-ms pulses of $\sim 100 \text{ mW per mm}^2$ intensity were delivered to the silicon probe insertion site using a MBL-III 473 nm 150 mW blue light laser (CNI, Chungchun, China). Reliable, short latency (<5 ms) and low jitter (<1 ms) responses were found in $\sim 5\%$ of the units, indicating that $\sim 95\%$ of cells were driven polysynaptically rather than directly.

Data analysis (electrophysiology). Spikes were detected and visually verified using the programs NDmanager and Neuroscope³⁵. Spike sorting involved an automated

stage, performed using KlustaKwik³⁶, and a manual verification stage for which Klusters³⁵ or KlustaView³⁷ were used. Detailed analysis of coupling to population rate and LFP was performed only for units with isolation distance >20 (see refs 38, 39). Units were selected and sorted blind to measures of population coupling and all other cellular parameters.

Population rate (for example, in Fig. 1c) was computed by accumulating all the detected spikes (both well-isolated units and multi-unit activity) with 1 ms resolution, and smoothing the resulting vector with a Gaussian of half-width 12 ms. The population rate used to compute the stPR for any individual unit did not include the spikes of that unit. The baseline level of each stPR (which reflects the mean population rate) was subtracted.

For stLFP computations, raw extracellular signals were first digitally band-pass filtered offline between 0.1 and 200 Hz to isolate the LFP. For units recorded on a particular shank, the LFP was taken from an adjacent shank (200 μm away), to avoid contaminating the stLFP by the spike waveform itself. The size of stLFP was taken as the ordinate value of the negative peak of the cross-correlation in a 1 s interval around 0 lag. stLFPs were normalized similarly to stPRs (see below).

The size of the stPR was quantified as the value of the spike-triggered population rate at 0 time lag. Thus, the population coupling of unit i is given by:

$$c_i = \frac{1}{\|f_i\|} \int f_i(t) \sum_{j \neq i} (f_j(t) - \mu_j) dt$$

Here, f represents the smoothed firing rate of a neuron (Gaussian kernel of half width $12/\sqrt{2}$ ms), μ is its mean firing rate, and $\|f\|$ represents its norm (that is, the number of spikes fired). To compare the sizes of stPRs across recordings, they were normalized by the median size of the stPR of the shuffled data in each recording (see next paragraph).

Spike shuffling was performed according to the previously described raster marginals model¹⁴. In more detail, the recording was first divided into non-overlapping 1 ms bins. A binary matrix was then constructed with one column for each time bin and one row for each isolated unit as well as additional rows for the multiunit spiking on each shank. Each matrix element contained a 1 if the corresponding unit spiked in the corresponding time bin. To shuffle, random 2-by-2 submatrices were repeatedly chosen with each row and column of the submatrix containing a 0 and 1; the positions of 0s and 1s were then exchanged in the submatrix, which leaves the summed values of each row and column identical. As we have discussed previously¹⁴, such a shuffling procedure produces in the limit a uniform sample from a distribution subject to the constraints on the mean firing rate and population rate distribution of the original data.

To characterize the responses of individual cells to drifting grating stimuli, the response for each orientation was averaged across trials, contrasts and spatial frequencies, and the orientation with the highest value was taken as preferred. The spatial frequency that evoked the highest response along the preferred orientation was taken as the preferred spatial frequency. The orientation selectivity index (OSI) was computed as: $(R_{\text{pref}} - R_{\text{ortho}})/(R_{\text{pref}} + R_{\text{ortho}})$, where R_{pref} and R_{ortho} are the responses in the preferred and orthogonal orientations. f_1/f_0 was taken as the ratio between the power of the average response around 2 Hz (which was the temporal frequency of the drifting grating stimuli) and the mean increase in the firing rate above the spontaneous level.

In primate area V4, we measured the firing rate changes during saccade preparation for isolated single neurons. Spikes were counted during the interval between 0.5 s after cue onset and the end of the post-cue period (that is, the start of the blank period), and were converted to firing rates for each trial based on the duration of that period. We compared firing rates from trials for which the cue indicated a saccade into the receptive field would be required and trials for which the cue indicated a saccade to an orthogonal location outside the receptive field would be required. Only spikes from correctly performed trials were considered for this analysis. Peak stPR size was measured identically to rodent electrophysiological recordings, and was computed from the continuous recording of the entire experimental session.

Two-photon imaging. All procedures were performed under licenses from the UK Home Office in accordance with the Animals (Scientific Procedures) Act 1986. In the imaging experiments C57BL/6J mice of P22–26 age and both sexes were used. Anaesthesia was induced with fentanyl, midazolam and medetomidine (0.05 mg kg^{-1} , 5.0 mg kg^{-1} and 0.5 mg kg^{-1} , correspondingly) and later maintained by isoflurane (0.3 – 0.5%) in a 60:40 mixture of oxygen and nitrous oxide. The experimental setup for two-photon imaging was described in refs 24, 25. In brief, visual stimuli were presented on an LCD monitor ~ 20 cm from the animal, covering a field of view of $\sim 105^\circ \times 85^\circ$. The calcium-sensitive dye Oregon Green BAPTA-1 (OGB-1) was bulk loaded into the superficial layers of the cortex together with sulforhodamine 101 to distinguish glia from neurons.

We used two data sets in the analysis, recorded from 18 and 15 animals, respectively. In the first data set (previously described in ref. 24), visual stimuli included

full-field and localized drifting gratings and natural movies. The gratings were shown at 100% contrast, 2 Hz temporal frequency, spatial frequency of 0.035 cycles per degree, and 8 directions. Cortical areas of $\sim 285 \times 285 \mu\text{m}$ were imaged at 7.6 Hz and 256×256 pixel resolution, and the imaging was performed for $\sim 8\text{--}16$ different depths, $7 \mu\text{m}$ apart.

In the second data set⁴⁰, static natural images and natural movies (either video clips from *The Life of Mammals* or cage scenes from a head-mounted mouse camera) were presented. A volume of $\sim 200 \times 200 \times 50 \mu\text{m}$ was imaged simultaneously at $512 \times 512 \times 4$ pixel resolution using a piezoelectric objective mover. In each experiment 2–3 such volumes were imaged.

After *in vivo* imaging, whole-cell *in vitro* recordings of the imaged neurons were performed as previously described²⁴. In brief, at the end of the *in vivo* imaging session the brain was removed, $300\text{-}\mu\text{m}$ thick coronal slices were cut and the slice containing the *in vivo* imaged region was located by the presence of OGB-1. A simultaneous recording from several nearby neurons was performed using a standard whole-cell *in vitro* protocol. Synaptic connections between neurons were tested using 30-Hz trains of five suprathreshold 5-ms current pulses, and checking for postsynaptic responses in the other cells. The current injection train was repeated at least 30 times with 15-s intervals. Pyramidal cells and interneurons were distinguished by spike width, resting membrane potential, responses to 1-s depolarizing currents, and somatodendritic morphology when filled with Alexa dye. For the analysis of the relationship between coupling strength and probability to receive (provide) a synaptic connection (Fig. 4), we excluded pairs where the potential target (source) cell was an interneuron or a neuron for which no *in vivo* calcium trace was available. Simultaneously patched neurons could have been imaged non-simultaneously, which happened when the cells resided in different imaging planes or volumes. In total, 379 neurons were recorded; of 854 potential input connections tested, 146 synaptic connections were found. Of 854 potential output connections tested, 155 synaptic connections were found. The number of potential input and output connections were equal by coincidence; some of the input connections were not analysed as output connections because one of the cells was either an interneuron or had no *in vivo* data, and vice versa.

Data analysis (imaging). Outlines of recorded neurons were detected semi-automatically using custom software written in MATLAB (MathWorks, Natick, Massachusetts). All pixels within individual outline were averaged to give for each neuron a single $\Delta F/F$ signal, which was additionally high-pass filtered above 0.02 Hz to correct for slow artefacts such as photobleaching.

The population coupling of each cell i was estimated by a similar formula to that used in spike train data:

$$c_i = \frac{1}{\|f_i\|} \int f_i(t) \sum_{j \neq i} (f_j(t) - \mu_j) dt,$$

where f now represents the continuous fluorescence trace of a cell ($\Delta F/F$), and $\|f\|$ its standard deviation. Note that c_i is proportional to the Pearson correlation of cell i with the summed activity of all other neurons. To pool data across multiple experiments, the normalization factor was computed not by spike shuffling (which is not possible for calcium traces) but by z-scoring the population coupling values within each recording.

For grating responses, the orientation selectivity index (OSI; Extended Data Fig. 6g) was estimated as $(R_{\text{best}} - R_{\text{ortho}})/(R_{\text{best}} + R_{\text{ortho}})$, where R_{best} and R_{ortho} are the interpolated responses to drifting gratings in the best direction and the directions orthogonal to it.

Statistics. Unless explicitly stated otherwise, comparisons between pairs of sets of values were made using Wilcoxon rank sum test and correlations were measured using Spearman's rank correlation coefficients (an important exception being the pairwise spike train correlations, for which, as commonly accepted, we used Pearson correlation). No statistical methods were used to predetermine sample size.

Pairwise correlation prediction. To predict pairwise correlations from coupling values, we used an extension of the raster marginals model¹⁴ that now also included a population coupling parameter for each neuron. The spike trains of all neurons are first converted into a binary 0–1 matrix with 20-ms bins. The parameters extracted correspond to: (i), the total number of 1 s in row i , for all i between 1 and the number of rows n ; (ii), the total number of columns whose sum is i , for all i between 0 and n ; and (iii) the inner product of each row with the sum of all rows (representing the sTPR for the unit corresponding to the row).

Pairwise correlations were predicted by constructing a random matrix that respects the constraints represented by the above parameters, and computing the correlations between all pairs of rows in this random matrix. To construct such a random

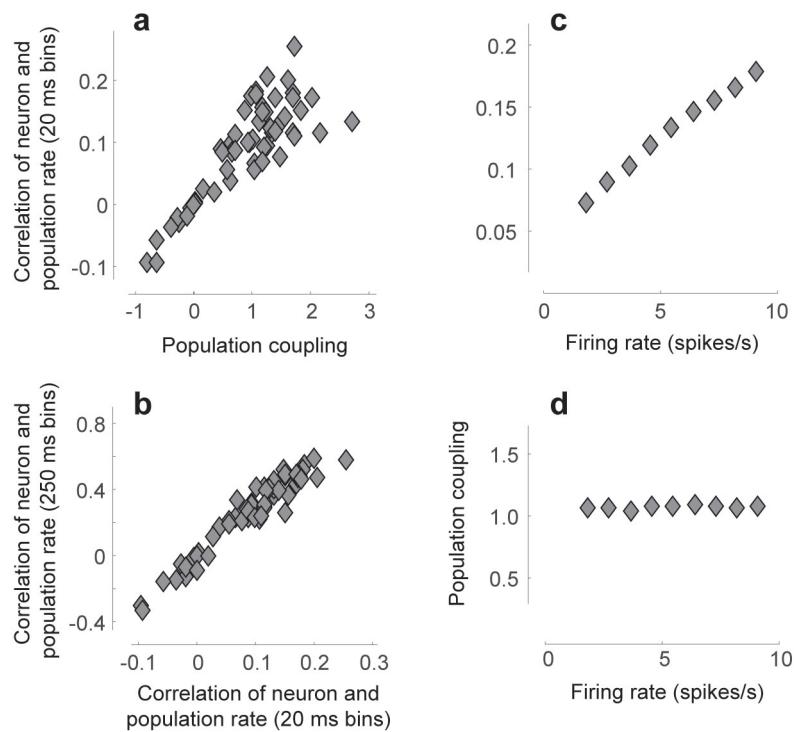
matrix we first construct a random matrix that respects constraints (i) and (ii), as described in ref. 14, and then perform a sequential operation to impose satisfaction of constraint (iii). This sequential operation consists of repeatedly exchanging 2×2 sub-matrices between rows whose population coupling is too high and too low. Specifically, observe that if there exists a row n_i for which $\left\langle n_i, \sum_j n_j \right\rangle$ (the inner product with the sum of all rows) is too high, then there must also exist another row for which this inner product is too low; this is because $\sum_i \left\langle n_i, \sum_j n_j \right\rangle = \left\langle \sum_i n_i, \sum_j n_j \right\rangle$, which depends only on the population rate distribution and is thus invariant to sub-matrix exchange.

The satisfaction of constraint (iii) occurs by repeatedly finding a pair of rows that violate the constraint in opposite directions, while such pairs exist. For each such pair of rows, we find a pair of columns that forms a 2-by-2 sub-matrix such that each row and column of it contains a 0 and 1, and switch the 0s and 1s in a way that preserves constraints (i) and (ii) and reduces the violation of constraint (iii) in this pair of rows. This procedure is repeated until (iii) is satisfied (up to a small error of at most n).

To measure how well this method could predict pairwise correlations, we used cross-validation. Time bins were divided at random into an equally-sized training and test set. Model parameters were fit from the training set, and the prediction error of the model was quantified as $SS_{\text{modelErr}} = \sum_{i < j} (c_{ij} - \hat{c}_{ij})^2$, where c_{ij} is the correlation between neurons i and j measured in the test set, and \hat{c}_{ij} is the correlation predicted by the model with parameters estimated on the training set. This error was compared to a 'ground truth' prediction, $SS_{\text{dataErr}} = \sum_{i < j} (c_{ij} - c'_{ij})^2$, where c'_{ij} is the correlation between neurons i and j as directly measured in the training set. The quality of the model fit was defined to be the fraction of explainable variance: $[SS_{\text{tot}} - SS_{\text{modelErr}}]/[SS_{\text{tot}} - SS_{\text{dataErr}}]$, where $SS_{\text{tot}} = \sum_{i < j} (c_{ij} - \bar{c})^2$ and \bar{c} is the mean pairwise correlation.

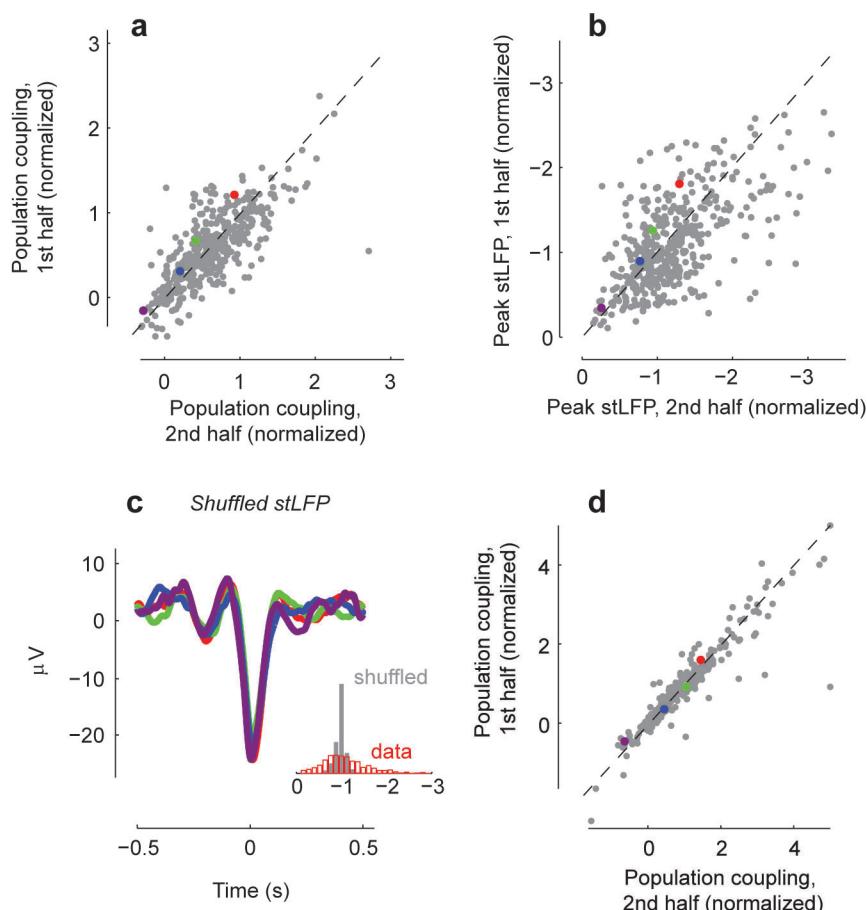
Latent variable analysis of the population recordings was performed using the publicly available toolbox of L. Buesing and J. Macke (http://bitbucket.org/mackelab/pop_spike_dyn; see also ref. 21). The toolbox uses a nuclear norm minimization method⁴¹ to fit a low-dimensional model of the form $s_t \sim \text{Poisson}(e^{y_t})$, where $y_t = Cx_t + d$. The latent variables x_t were modelled using a multivariate Gaussian distribution. The performance of this model was assessed by its ability to predict pairwise correlations, quantified with cross-validation measures described above. To generate correlations we fitted the parameters of the model and then used these parameters to sample spike count vectors.

31. Margrie, T. W., Brecht, M. & Sakmann, B. *In vivo*, low-resistance, whole-cell recordings from neurons in the anaesthetized and awake mammalian brain. *Flieggers Arch.* **444**, 491–498 (2002).
32. Haider, B., Hausser, M. & Carandini, M. Inhibition dominates sensory responses in the awake cortex. *Nature* **493**, 97–100 (2013).
33. Ayaz, A., Saleem, A. B., Scholvinck, M. L. & Carandini, M. Locomotion controls spatial integration in mouse visual cortex. *Curr. Biol.* **23**, 890–894 (2013).
34. Luczak, A., Bartho, P., Marguet, S. L., Buzsaki, G. & Harris, K. D. Sequential structure of neocortical spontaneous activity *in vivo*. *Proc. Natl Acad. Sci. USA* **104**, 347–352 (2007).
35. Hazan, L., Zugaro, M. & Buzsaki, G. Klusters, NeuroScope, NDManager: a free software suite for neurophysiological data processing and visualization. *J. Neurosci. Methods* **155**, 207–216 (2006).
36. Harris, K. D., Henze, D. A., Csicsvari, J., Hirase, H. & Buzsaki, G. Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J. Neurophysiol.* **84**, 401–414 (2000).
37. Rossant, C. & Harris, K. D. Hardware-accelerated interactive data visualization for neuroscience in Python. *Front. Neuroinform.* **7**, 36 (2013).
38. Harris, K. D., Hirase, H., Leinekugel, X., Henze, D. A. & Buzsaki, G. Temporal interaction between single spikes and complex spike bursts in hippocampal pyramidal cells. *Neuron* **32**, 141–149 (2001).
39. Schmitzer-Torbert, N., Jackson, J., Henze, D., Harris, K. & Redish, A. D. Quantitative measures of cluster quality for use in extracellular recordings. *Neuroscience* **131**, 1–11 (2005).
40. Cossell, L. *et al.* Functional organization of excitatory synaptic strength in primary visual cortex. *Nature* **518**, 399–403 (2015).
41. Dorn, J. D. & Ringach, D. L. Estimating membrane voltage correlations from extracellular spike trains. *J. Neurophysiol.* **89**, 2271–2278 (2003).
42. de la Rocha, J., Doiron, B., Shea-Brown, E., Josic, K. & Reyes, A. Correlation between neural spike trains increases with firing rate. *Nature* **448**, 802–806 (2007).
43. Niell, C. M. & Stryker, M. P. Highly selective receptive fields in mouse visual cortex. *J. Neurosci.* **28**, 7520–7536 (2008).



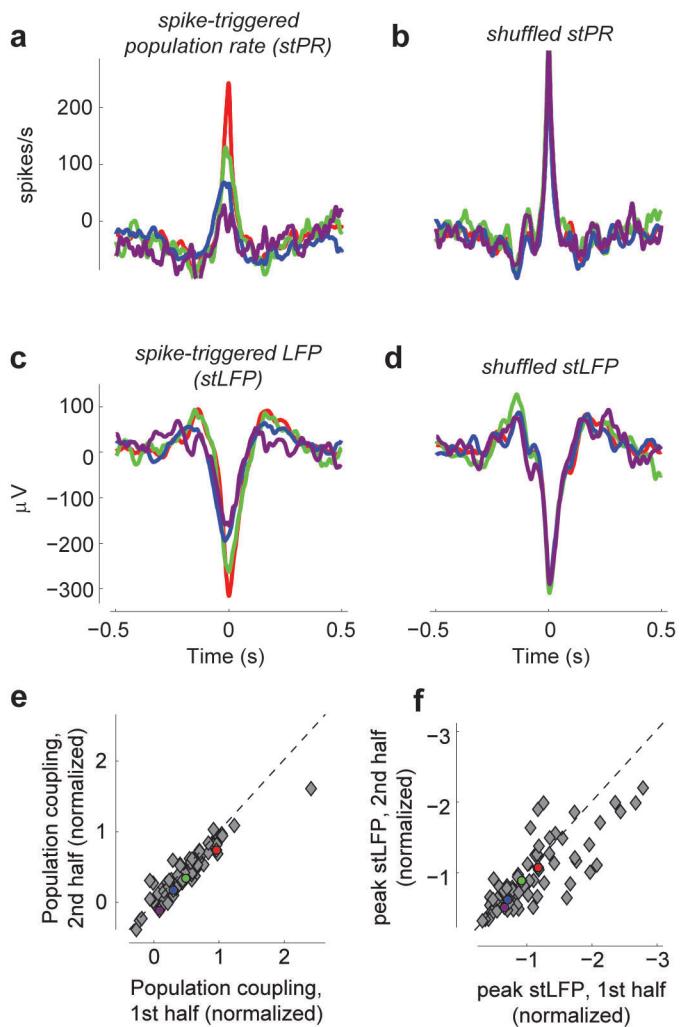
Extended Data Figure 1 | Pearson correlation between spike trains of individual units and the population rate. To estimate the relation of a neuron to the population, an alternative to spike-triggered population rate (stPR) would have been to compute the Pearson correlation coefficient of the neuron's spike train with the summed population rate of all other recorded cells (a measure we term 'Pearson coupling'). This measure, however, is biased by firing rate. **a**, Pearson coupling and stPR were computed for a set of individual units in an example experiment. Pearson coupling is related to the stPR, but not identical to it. **b**, The numerical value of the Pearson coupling depends

strongly on the bin size used, but the correlations measured with different bin sizes are tightly related. **c**, Pearson correlation is biased by firing rate^{41,42}. The spike train of a single cell was 'thinned' to different firing rates by keeping only a random subset of its spikes; Pearson correlation with the population was recalculated for different values of firing rate. A strong effect of firing rate is seen. **d**, Performing the same analysis for population coupling (measured by stPR) demonstrates that this measure does not suffer from rate bias. For this reason, we chose to quantify population coupling with stPR in this work.

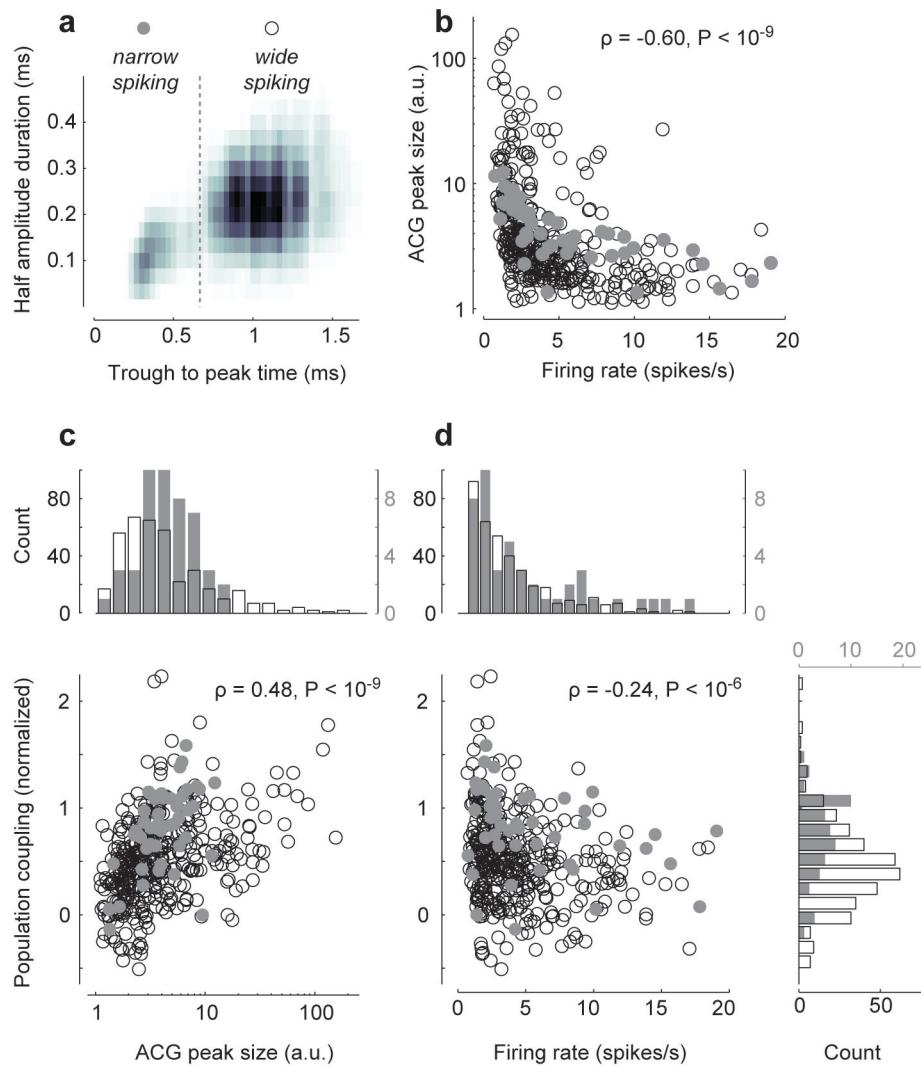


Extended Data Figure 2 | Neighbouring neurons differ markedly in population coupling during spontaneous activity. **a**, Dividing the data into two halves shows that population coupling, measured as the height of stPR at 0 time lag, is highly consistent over time ($n = 431$ neurons from 13 experiments; $\rho = 0.76$, $P < 10^{-100}$, rank correlation). Coloured dots represent the four example cells. **b**, As in **a** for peak spike-triggered local field potential (stLFP) ($\rho = 0.58$, $P < 10^{-100}$). **c**, Differences in stLFP disappear after

shuffling spikes in a manner that preserves each neuron's mean firing rate and the population rate (compare with Figure 1g). Inset, stLFPs in the actual spike trains (red) and after shuffling (grey), for neurons from all experiments (compare with Figure 1g). **d**, stPR size of V4 neurons is consistent over time ($n = 262$ neurons from 43 experiments; $\rho = 0.95$, $P < 10^{-100}$, rank correlation).

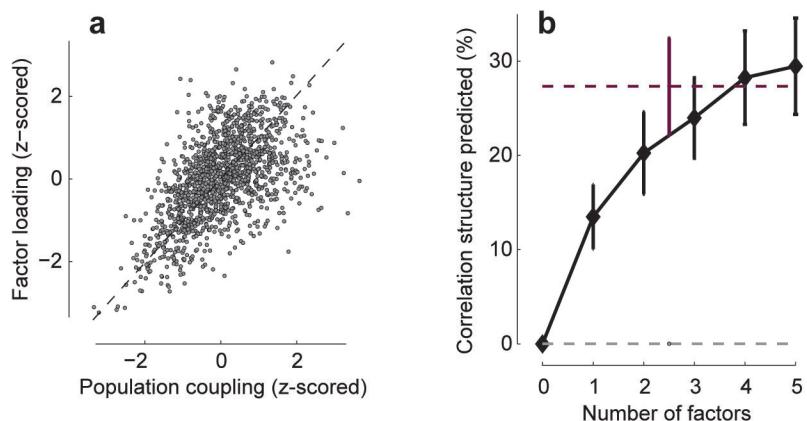


Extended Data Figure 3 | Neighbouring neurons in auditory cortex differ markedly in population coupling. **a**, Spike-triggered population rate (stPR) for four example neurons recorded on the same electrode shank, during spontaneous activity in rat primary auditory cortex. **b**, Differences in population coupling disappear after shuffling spikes in a manner that preserves each neuron's mean firing rate and the population rate distribution. **c, d**, As in **a, b** for the spike-triggered local field potential (stLFP). **e**, Dividing the data into two halves shows that population coupling, measured as the height of stPR at 0 time lag, is highly consistent over time ($n = 76$ neurons from 3 experiments; $\rho = 0.92$, $P < 10^{-100}$, rank correlation). Coloured dots represent the four example cells. **f**, As in **e** for stLFP ($\rho = 0.81$, $P < 10^{-100}$, rank correlation).



Extended Data Figure 4 | Firing rate, burstiness and population coupling.
a, Similarly to other studies^{20,43} our recordings allow separation of narrow spiking (putative *Pvalb*⁺ inhibitory) and wide spiking (primarily excitatory pyramidal) neurons. Here, we used a trough-to-peak time of 0.66 ms as the separation criterion. **b**, There is a negative correlation between burstiness (the ratio between the peak and baseline of a neuron's autocorrelogram) and mean firing rate, which is also the case individually for wide spiking ($n = 384, \rho = -0.60, P < 10^{-9}$, rank correlation) and narrow spiking ($n = 47, \rho = -0.82, P < 10^{-9}$, rank correlation) neurons. a.u., arbitrary units. **c**, There is a positive correlation between burstiness and population coupling, which is also the case individually for wide spiking ($\rho = 0.46, P < 10^{-9}$, rank correlation)

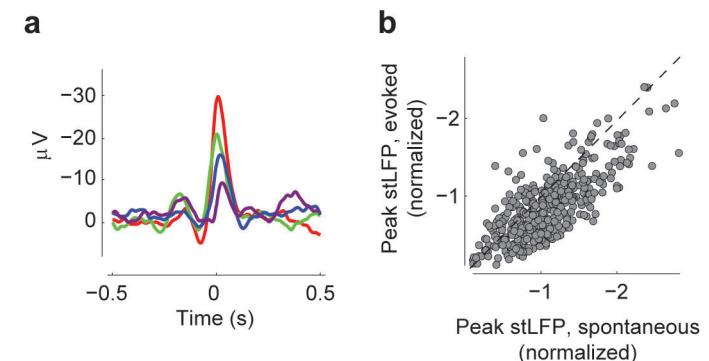
and narrow spiking ($\rho = 0.50, P = 4 \times 10^{-4}$, rank correlation) neurons. **d**, There is a negative correlation between firing rate and population coupling, which is also the case individually for wide spiking ($\rho = -0.27, P = 10^{-7}$, rank correlation) and narrow spiking ($\rho = -0.37, P = 0.01$, rank correlation) neurons. The correlation between population coupling and firing rate can be predicted from the correlations between burstiness and firing rate and between population coupling and burstiness; the partial rank sum correlation between population coupling and firing rate, once burstiness is taken into account, is insignificant ($\rho = 0.06, P = 0.25$). This is also the case for wide spiking ($\rho = 0.01, P = 0.78$) and narrow spiking ($\rho = 0.07, P = 0.65$) neurons individually.



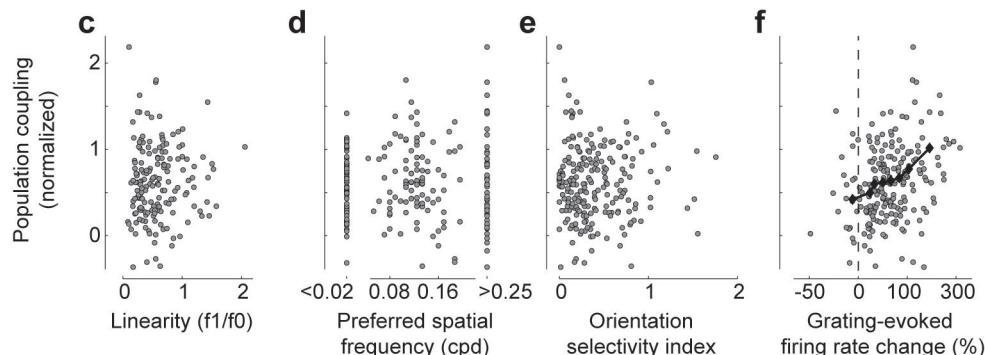
Extended Data Figure 5 | Latent variable analysis. **a**, Population coupling of each neuron is highly correlated with its loading in a single-factor latent variable model (see Methods). The similarity of each cell's population coupling and loading indicates that the low-dimensional structure found by the latent variable model is homologous to that found by the coupling model. **b**, Percent of pairwise correlation structure explained by a latent variable model with 1–5 factors (black), and by the coupling model introduced in the present study

(dashed purple line). Error bars show standard error. While the coupling model outperforms latent variable models with less than four degrees of freedom, this difference may arise primarily from the assumption of a Gaussian distribution for the latent variables. Indeed, if the population rate distribution generated by the latent variable model is substituted into the coupling model instead of the (correct) population rate distribution, extremely poor performance results (dashed grey line).

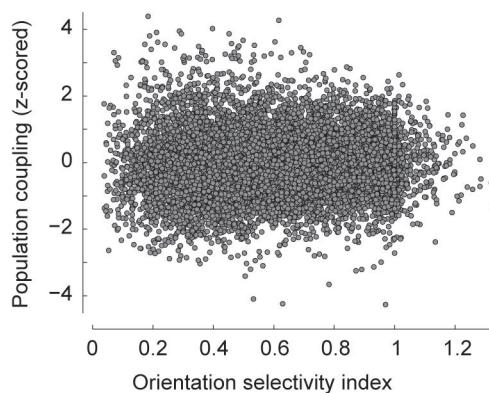
Electrophysiological data - natural movies



Electrophysiological data - drifting gratings

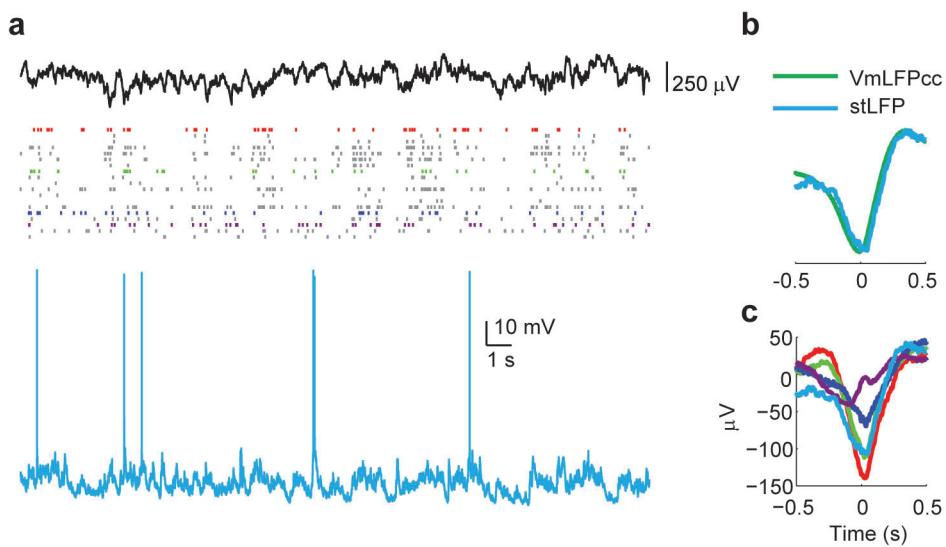


g Imaging data - drifting gratings



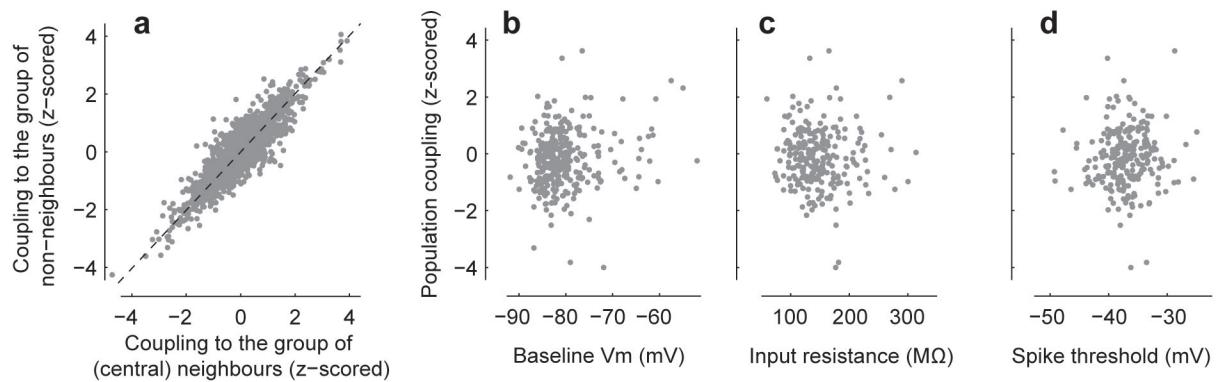
Extended Data Figure 6 | Population coupling and visual stimulation in mouse V1. **a**, stLFPs computed for the four example neurons of Fig. 1a–f, from intervals of natural movie presentation (inverted for ease of comparison, see Figure 1f). **b**, Comparison of stLFP size during spontaneous and evoked activity across all experiments ($\rho = 0.72, P < 10^{-100}$, rank correlation). **c–e**, Population coupling is plotted versus the f_1/f_0 ratio, preferred spatial frequency and orientation selectivity index (OSI) for neurons recorded in the infragranular layers of V1. All correlations are statistically insignificant. **f**, Similar to movie

presentations (Figure 3e), the mean change in the activity of a cell in response to grating presentations (relative to baseline, averaged across contrasts and orientations) correlates with population coupling measured during spontaneous activity ($\rho = 0.32, P = 2 \times 10^{-6}, n = 217$, rank correlation). Black diamonds, running median. **g**, In the two-photon imaging data (of $\sim 10,000$ cells) only a very weak correlation between OSI and population coupling was found ($\rho = 0.066, P < 10^{-9}$, rank correlation).



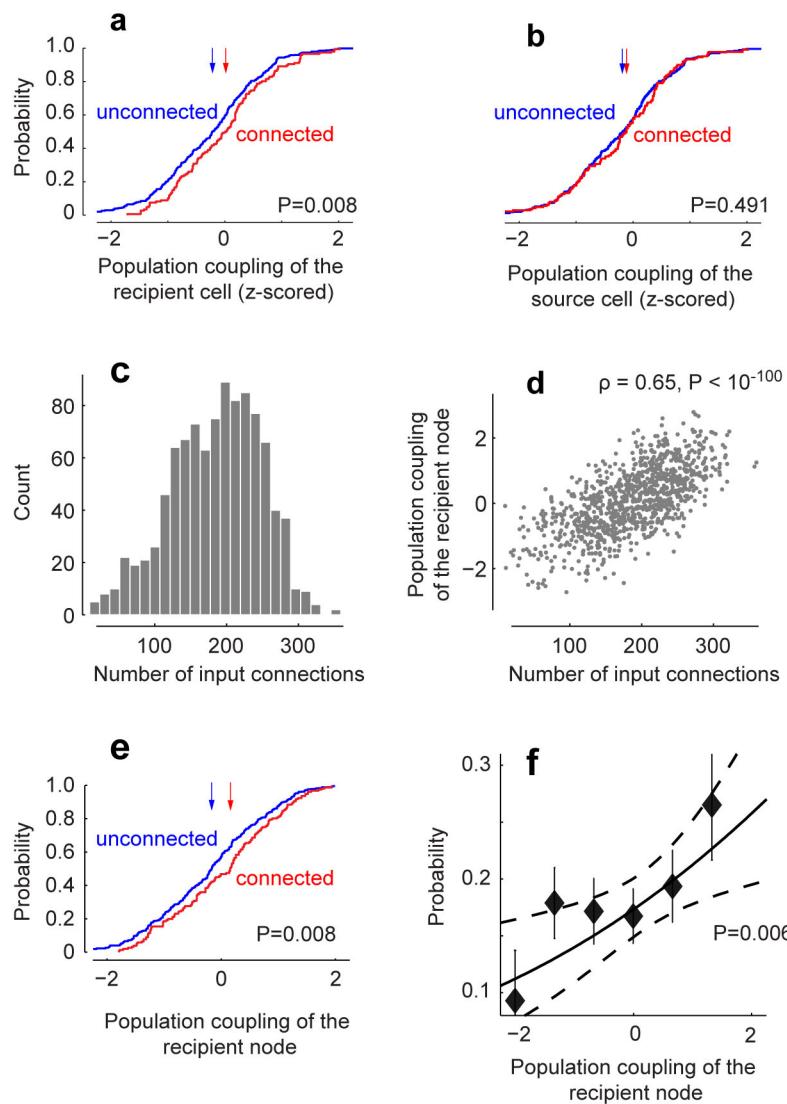
Extended Data Figure 7 | stLFP reflects the correlation between membrane potential (V_m) and LFP. **a**, Example of a silicon probe population recording performed simultaneously with a whole-cell recording (in an anaesthetized animal). Four neurons shown in colour were recorded on the same shank of the

silicon probe. **b**, Comparison of stLFP and V_m -LFP cross-correlation ($VmLFPcc$, appropriately scaled along the ordinate axis) for the intracellularly recorded cell. **c**, stLFP for the four neurons from **a** and the intracellularly recorded neuron, exhibiting diversity in the strength of coupling to LFP.



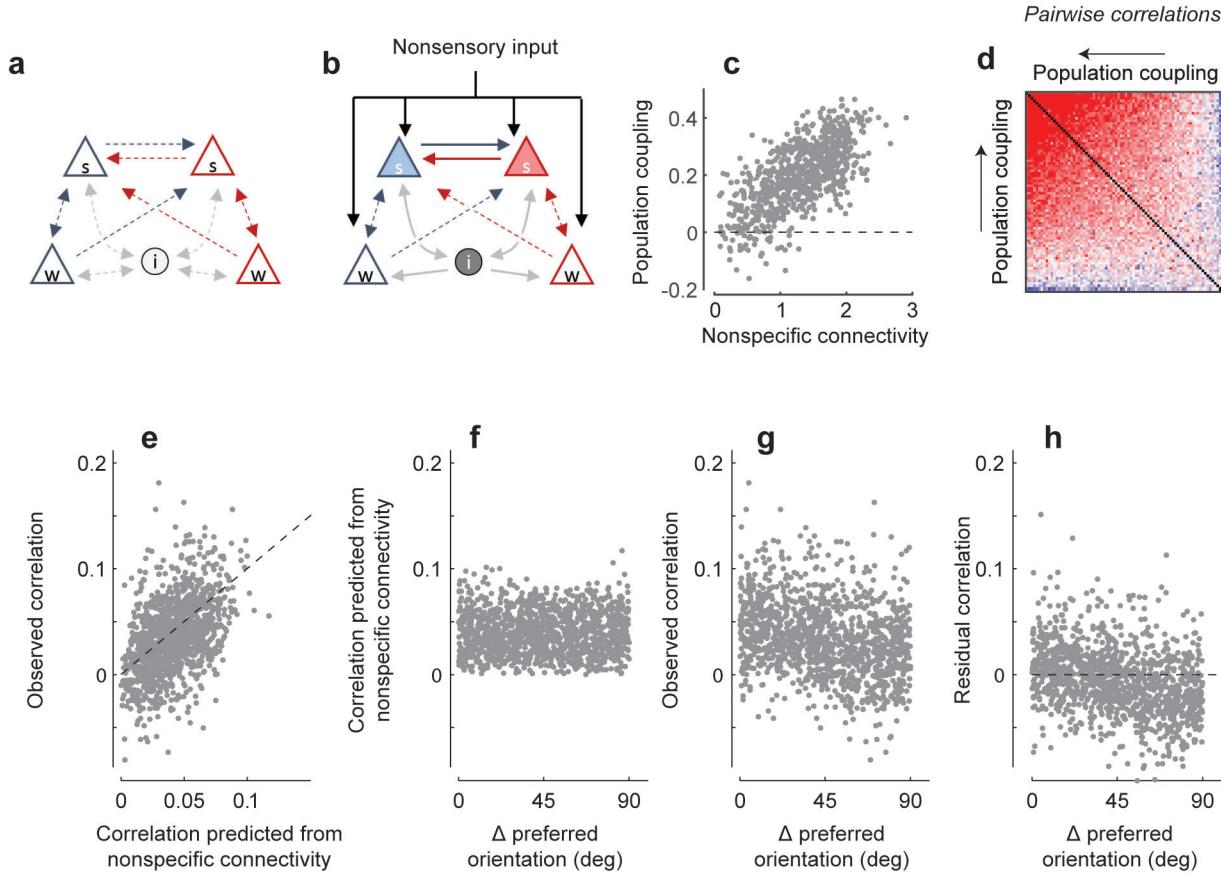
Extended Data Figure 8 | Population coupling in two-photon data is not correlated with location and intrinsic properties of the neurons. **a**, For each neuron in the central region of the imaging field (defined as a square quarter of the total imaging area), we compared its coupling to the population of all other neurons in the central region, with its coupling to population of all neurons outside of the central region. The two were highly similar; this was the case because the two population rate signals were themselves highly correlated

(on average across experiments the Pearson correlation was 0.77). Thus, differences in population coupling measured between cells do not reflect differences in the fraction of nearby neurons imaged. **b–d**, No significant correlation was observed between population coupling (measured *in vivo*) and resting potential, input resistance and spike threshold (subsequently measured *in vitro*).



Extended Data Figure 9 | Correlation between input connectivity and population coupling. **a**, Cumulative distribution of population coupling of a target pyramidal neuron when an input connection was present (red) and when it was absent (blue). The medians (arrows) are significantly different ($P = 0.008$, rank sum test, $n = 854$ pairs). **b**, As in **a** for population coupling of the source pyramidal cells. The distributions shown were used for the logistic regression analysis in Fig. 4. **c–f**, To estimate what strength of correlation between input connectivity and population coupling would give rise to these observations, we constructed random directed graphs of 1,000 nodes (each node representing a L2/3 pyramidal cell) with the probability of connection from node j to node i given by $\sqrt{p_{in}^i p_{out}^j}$, where the propensities to receive and provide connections (p_{in} and p_{out} , correspondingly) were randomly and independently chosen for each node from a Gaussian distribution. The resulting distribution of the number of input connections in a typical network is shown in **c**; the number of output connections was (by construction) similarly distributed. In addition, each node was assigned a population coupling

value, highly correlated to the number of its input connections (on average $\rho = 0.65$); this correlation in a typical network is shown in **d**. **e, f**, We next asked how the relationship between measured connectivity and population coupling would look if we sample from 33 such randomly generated networks (equal to the number of animals used in our experimental data), the same amount of data empirically available in our *in vitro* recordings (that is, the connections between 2–3 randomly selected groups of 2–6 nodes). Results very similar to those of Fig. 4 were typically obtained (**e, f**; compare with **a** and Fig. 4d; error bars in **f** indicate standard error for binned data). In particular, when the entire procedure was repeated 1,000 times, in over 30% of the cases the P value of the difference between the medians (presented in **a, e**) was higher (that is, less statistically significant) than the value of 0.008 obtained in the actual data. Thus, the results shown in Fig. 4 and in **a** are consistent with a strong correlation between connection probability and population coupling.



Extended Data Figure 10 | Mathematical model for the relationship between nonspecific connectivity, specific connectivity, and correlations.

a, A recurrent network where excitatory cells (triangles) send synaptic connections (arrows) to each other and to inhibitory cells (circles). Weakly coupled neurons (bottom) receive only connections from neurons with similar sensory preference (for example, for stimulus orientation, indicated in blue versus red). Strongly coupled neurons (top) also receive nonspecific connections from neurons of different sensory preference. **b**, The effect of nonspecific drive, such as caused by non-sensory top-down inputs, or occurring due to artificial optogenetic stimulation, is amplified through recurrent connections, leading to stronger activation of neurons with greater mean local input (darker shading). **c, d**, Correlations predicted by the model (analytically derived in Supplementary Information). **c**, Population coupling versus nonspecific connectivity c_i , for all simulated excitatory neurons.

d, Pseudocolour plot of predicted pairwise correlations for a random subset of excitatory neurons, ordered by population coupling. **e–h**, Dependence of correlations on specific and nonspecific connectivity. **e**, Predicted correlations based on nonspecific connections versus total observed correlations. **f**, Predicted correlation based on nonspecific connectivity versus difference in preferred orientation. As in the experimental data (Fig. 2e), no relation is observed. **g**, Observed correlation versus difference in preferred orientation. As has been widely reported, observed correlations are largest for neurons of similar orientation preference. **h**, Residual correlation (after removing prediction from nonspecific connectivity) versus difference in preferred orientation. Again as in our experimental data (Fig. 2f), the residual correlation is largest for neurons of similar orientation preference, indicating an additive relationship between correlations generated by specific connections and correlations generated by nonspecific connections.

Counteraction of antibiotic production and degradation stabilizes microbial communities

Eric D. Kelsic¹, Jeffrey Zhao¹, Kalin Vetsigian² & Roy Kishony^{1,3}

A major challenge in theoretical ecology is understanding how natural microbial communities support species diversity^{1–8}, and in particular how antibiotic-producing, -sensitive and -resistant species coexist^{9–15}. While cyclic ‘rock–paper–scissors’ interactions can stabilize communities in spatial environments^{9–11}, coexistence in unstructured environments remains unexplained^{12,16}. Here, using simulations and analytical models, we show that the opposing actions of antibiotic production and degradation enable coexistence even in well-mixed environments. Coexistence depends on three-way interactions in which an antibiotic-degrading species attenuates the inhibitory interactions between two other species. These interactions enable coexistence that is robust to substantial differences in inherent species growth rates and to invasion by ‘cheating’ species that cease to produce or degrade antibiotics. At least two antibiotics are required for stability, with greater numbers of antibiotics enabling more complex communities and diverse dynamic behaviours ranging from stable fixed points to limit cycles and chaos. Together, these results show how multi-species antibiotic interactions can generate ecological stability in both spatially structured and mixed microbial communities, suggesting strategies for engineering synthetic ecosystems and highlighting the importance of toxin production and degradation for microbial biodiversity.

Antibiotic-producing species are common in natural microbial communities^{13,17,18}. Ecological models of antibiotics typically assume pairwise species relationships, where antibiotic producers inhibit sensitive species more than resistant species. These pairwise inhibitory interactions, combined with costs for production and resistance, can lead to relationships of cyclic dominance among species (for example, rock–paper–scissors games), which can support coexistence in spatial environments^{9–11} beyond the limit set by competitive exclusion^{4,5}. However, such pairwise interaction models lead to coexistence through the separation of species into distinct spatial domains, whereas in nature antibiotic-producing, -resistant and -sensitive species appear to intermix even at very small length scales^{17,19}. Furthermore, species communities stabilized through pairwise antibiotic interactions are not resilient to the high level of species dispersal expected in nature^{12,16}. Understanding how multiple antibiotic-producing species coexist despite dispersal remains an open question.

The inhibitory interaction between an antibiotic-producing species and an antibiotic-sensitive species can be attenuated by the presence of a third ‘modulator’ species (Fig. 1a). One established mechanism for antibiotic attenuation is enzymatic degradation²⁰, a common mechanism for antibiotic resistance²¹. In principle, a modulator species could also intensify inhibitory interactions between two species^{22,23}, for example by inducing antibiotic production. However, when testing for such interactions among a collection of soil isolates using a three-species interaction assay (Fig. 1b), we observed that intensification was rare, while attenuation was common (Extended Data Figs 1, 2 and Extended Data Table 1). Realizing that such three-way attenuating interactions commonly occur among natural species motivated us to

explore their impact on ecological dynamics. We focus on antibiotic attenuation caused by degradation, which we observed experimentally (Extended Data Figs 1d and 2c), but our analysis can be generalized to other attenuation mechanisms such as antibiotic suppression²⁴. It is known that antibiotic-degrading species can coexist together with sensitive species when an antibiotic is provided externally^{25–27}. However, when antibiotics are produced by the species themselves, these two-species communities are no longer stabilized by degradation^{27,28}. The impact of antibiotic-degrading species on the stability of larger ecosystems has not been explored.

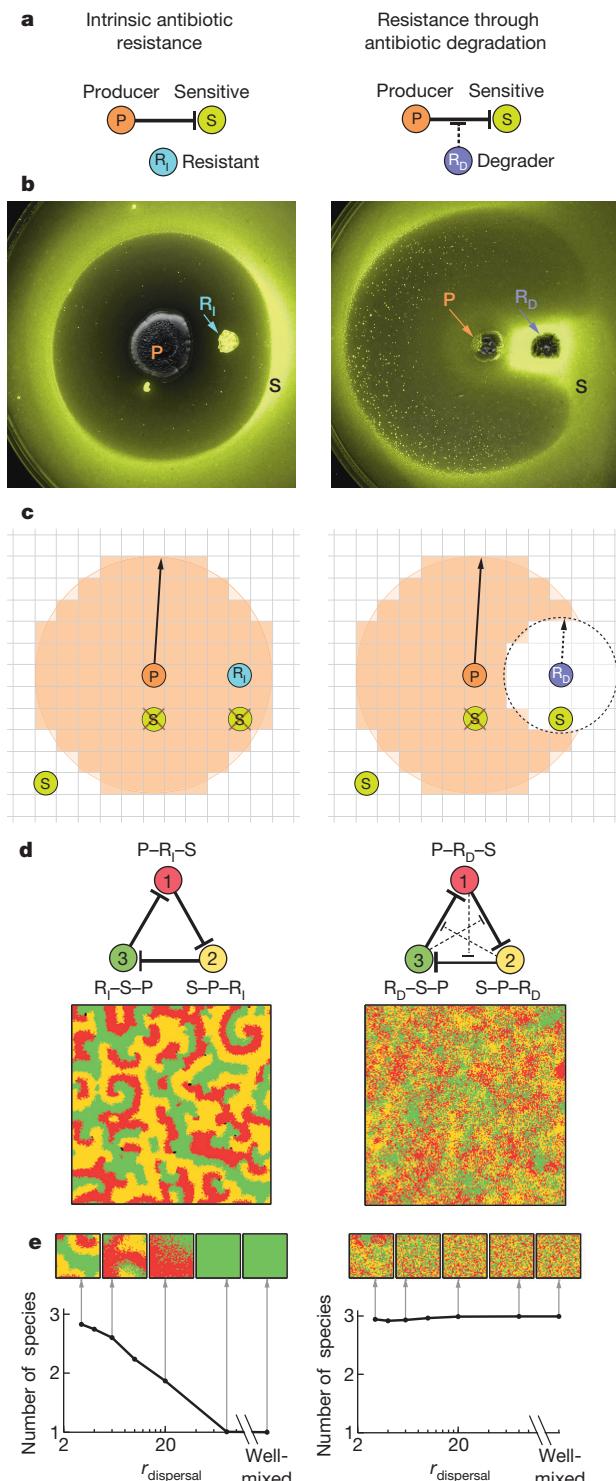
We investigated how antibiotic-degrading species affect the dynamics of microbial communities containing antibiotic producers by modifying a classical spatial model of antibiotic-mediated interactions. As in previous spatial models of antibiotic inhibition^{9–12}, we consider antibiotic-producing, -sensitive or -resistant species phenotypes. However, unlike in previous models, in addition to intrinsic resistance we also consider resistance through antibiotic degradation (Fig. 1a, b). Antibiotic-degrading species remove antibiotics from nearby locations, thereby protecting not only themselves but also neighbouring species. The simulations are performed on a grid (Fig. 1c), with each simulation step consisting of: production of antibiotics around antibiotic-producing species (within area of size K_p), removal of antibiotics near antibiotic-degrading species (within area of size K_d), killing of antibiotic-sensitive species within the antibiotic zones, and finally colonization of empty regions on a new grid by randomly choosing surviving species within a given dispersal radius, $r_{\text{dispersal}}$ (Extended Data Fig. 3; Methods).

In these spatial inhibition-zone models, communities with intrinsic resistance and with resistance through antibiotic degradation result in dramatically different patterning and robustness to dispersal. We simulated a simple three-species and three-antibiotic network that exhibits cyclic dominance. Consistent with previous studies^{9–12}, pairwise interactions among intrinsically resistant species result in coexistence of all three species through single-species domains that continually chase each other around the grid (Fig. 1d, left). However, we found that three-way interactions created by antibiotic degradation lead to tight intermixing of species (Fig. 1d, right). This fine-scale intermixing allowed us to investigate whether a spatial environment was necessary for the coexistence of these antibiotic-degrading communities. Spatial structure is needed for coexistence in pairwise antibiotic-interaction models^{9–12}, and diversity of the community with intrinsic resistance collapses when the dispersal radius increases (Fig. 1e, left; Supplementary Videos 1–5). In contrast, the antibiotic-degrading community maintains diversity across any level of dispersal, even with complete mixing between time steps (Fig. 1e, right; Supplementary Videos 6–10).

We characterized the stability of these communities with respect to model parameters by considering the analytical limit of an infinitely large environment with complete mixing between time steps ($r_{\text{dispersal}} \rightarrow \infty$; a ‘mixed inhibition-zone model’; Methods). Considering different initial species abundances, the community with

¹Department of Systems Biology, Harvard Medical School, Boston, Massachusetts 02115, USA. ²Department of Bacteriology and Wisconsin Institute for Discovery, University of Wisconsin-Madison, Madison, Wisconsin 53715, USA. ³Department of Biology and Department of Computer Science, Technion-Israel Institute of Technology, Haifa 3200003, Israel.

intrinsically resistant species coexists only at a single unstable fixed point; starting in any other initial condition leads to extinction of all but one species (Fig. 2a, left). In contrast, the community with antibiotic degradation coexists with a stable fixed point and a large basin of attraction (Fig. 2a, right). Moreover, antibiotic degradation enables stable communities to form despite large differences in inherent species growth rates (g_1, g_2, g_3 ; Fig. 2b, right). Thus, even in environments with complete mixing, the interplay of antibiotic production and degradation leads to coexistence that is robust to large perturbations of species abundances and substantial differences in inherent growth rates.



Community stability is maximized at intermediate levels of antibiotic degradation (Fig. 2c and Extended Data Fig. 4a). Levels of degradation that are too high mostly eliminate the effects of the antibiotics, after which inherent growth rate differences lead to one species taking over. Levels of degradation that are too low do not allow coexistence, as expected, since $K_D = 0$ is equivalent to intrinsic antibiotic resistance. Stability is therefore maximal at intermediate levels of degradation where a negative feedback can operate, such that increased abundance of a given species results in its increased inhibition (Extended Data Fig. 4b-d). In contrast, stability is maximal at high levels of antibiotic production (K_P). Thus, strong antibiotic production can stabilize diversity when combined with intermediate levels of antibiotic degradation.

Production and degradation of multiple antibiotics also leads to robust coexistence in a well-mixed chemostat setting. We modelled the cyclic three-species, three-antibiotic interaction network using a single resource chemostat model. While these models are inherently unstable with intrinsic resistance, implementing resistance through enzymatic degradation of antibiotics allows coexistence of all species through stable fixed points or limit cycles (Methods). As in the inhibition-zone model, stability occurs despite inherent growth rate differences among species and was strengthened at high levels of production and intermediate levels of degradation (Extended Data Fig. 5a). We also observed robust coexistence for different forms of antibiotic inhibition functions and when inhibitory interactions are diffuse²⁹, that is, when all species have finite sensitivity to each antibiotic (Extended Data Fig. 5b). This generality with respect to different models, assumptions, and parameter values suggests that the counteraction of antibiotic production and degradation can support coexistence in a wide variety of settings.

These three-species ecosystems are also stable with respect to cheaters; that is, they resist invasion by species that gain a small growth advantage by ceasing production or degradation of antibiotics. We consider two types of cheating species: those that stop producing antibiotics and revert to intrinsic resistance while enjoying the production of antibiotics by their parental species, and those that stop degrading antibiotics to become antibiotic sensitive, yet still enjoy protection by their parental species (Fig. 3a). Starting with a stable three-species community, we simulated potential cheater invasion by adding a small amount of a cheater species derived from one of the existing species. We calculated the final abundance of a cheater once the community reached steady state while varying the growth advantage of the cheater compared to its parent: $g_i^{\text{cheater}} = (1 + \varepsilon)g_i$. We assume that each species has a different inherent growth rate, $g_1 < g_2 < g_3$. For small

Figure 1 | Replacing intrinsic antibiotic resistance with degradation-based resistance generates community robustness to species dispersal. **a**, Pairwise interactions among antibiotic-producing (P), -sensitive (S) and intrinsically resistant (R_I) species are compared to three-way interactions where the sensitive species can be protected by an antibiotic degrading species (R_D). **b**, Images of a YFP-labelled probe *Escherichia coli* strain (S, yellow-green) growing in the presence of two *Streptomyces* colonies. The inhibition of the probe strain by a producer (dark area around P) is unaffected by an intrinsically resistant species (R_I , left), but is strongly attenuated around an antibiotic-degrading species (right, yellow-green halo around R_D colony). Images are representative of cases without and with three-way interactions among 54 *Streptomyces* pairs tested (1 replica, Extended Data Figs 1 and 2). **c**, Spatial inhibition-zone model. A producer (P) kills nearby species that are sensitive to its antibiotic (crossed out S), but does not affect resistant species (R_I, R_D). Sensitive species are protected by degrader species (right, S within dashed circle around R_D). Surviving species then replicate and disperse over distance $r_{\text{dispersal}}$. **d**, Snapshots of spatial simulations for cyclical three-species, three-antibiotic interaction networks (labels indicate species phenotypes for each antibiotic). **e**, Intrinsically resistant communities collapse with increased dispersal (left), while communities with antibiotic degradation are robust to any level of dispersal (right). Insets show typical sub-region snapshots. Number of species is based on average Shannon diversity of sub-regions (Methods). $r_{\text{dispersal}}$ is measured in grid units (where individual grid cells have side length 1).

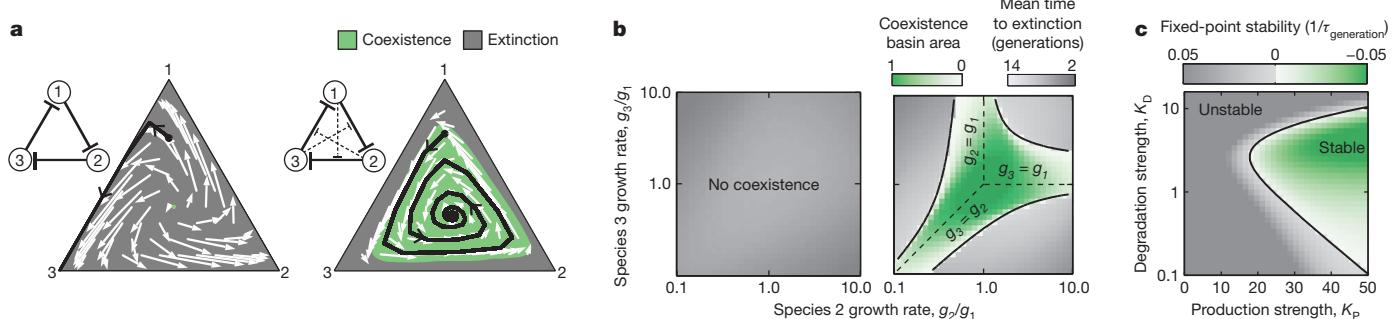


Figure 2 | Strong antibiotic production with intermediate levels of degradation leads to stable communities robust to initial conditions and substantial differences in species growth rates. **a**, Communities with intrinsic resistance diverge from an unstable fixed point (green dot, left), while communities with antibiotic degradation converge towards a stable fixed point within a large basin of attraction (green area, right). Shown in trilinear coordinates are abundance changes per generation (white arrows) and example trajectories (black) for species with inherent growth rates: $g_1 < g_2 < g_3$. **b**, Communities with degradation remain stable even for large differences in

ε , the only cheater that can invade the community is the production cheater derived from the fastest growing strain (Fig. 3b). This cheater replaces its parent, forming a new community of three species interacting through only two antibiotics. Coexistence in this simpler community is robust to parameter variation (Extended Data Fig. 6) and is also evolutionarily stable: it does not allow invasion by cheaters arising from any of its species (Extended Data Fig. 7). All other cheaters in the

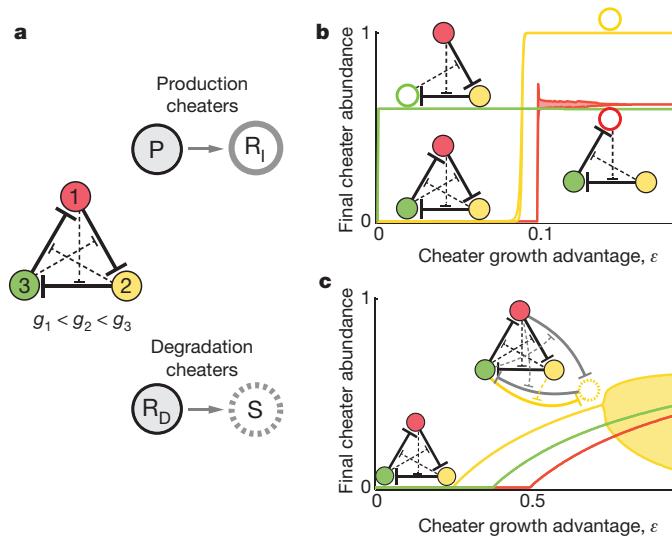


Figure 3 | Communities are robust to invasion by cheaters that cease antibiotic production or degradation. **a**, Production or degradation cheaters (open circles with solid or dashed lines, respectively) are introduced at low abundance into a stable three-species community (filled coloured circles). **b**, When cheater growth advantage ε is small, only the production cheater of the fastest growing species can invade (species 3, green). This cheater replaces its parent, forming a new stable community, which is in turn robust to further invasions. The other two production cheaters can invade only when ε is large enough, taking over the community (yellow) or replacing their parent in a stable community (red). **c**, Degradation cheaters cannot invade the community at low ε , whereas at high enough ε they can invade without replacing their parents, rather generating a stable four-species community. A four-species interaction network is shown for the species 2 degradation cheater: grey interactions are inherited from the parent species; yellow interactions are specific to the cheater. Shaded areas span lower/upper bounds when the community exhibits sustained oscillations.

inherent species growth rates (right), while communities with intrinsic resistance are always unstable (left). **c**, Maximal community stability occurs at high levels of antibiotic production (K_P) and intermediate levels of degradation (K_D). Shading shows rate of exponential convergence (green) or escape (grey) from the fixed point. Black lines in panels **b** and **c** show the stability region boundary based on linear stability analysis (Methods). Units for production and degradation strength are species $^{-1}$; $\tau_{\text{generation}}$ is the number of generations taken to get closer/further to the fixed point by a factor of 10.

three-species, three-antibiotic community cannot invade, even with a small growth advantage (that is, for small positive ε their final abundance is zero; Fig. 3c). With higher growth advantages, these cheaters can ultimately invade by dominating the community, by replacing the parent in a newly formed three-species community, or by generating a new four-species community (Fig. 3c). The emergence of these four-species communities motivated us to search for more complex ecosystems with greater numbers of species and to explore their dynamic behaviours.

Analysis of more complex networks identified many ways for antibiotic production and degradation to generate stability. Randomly sampling larger networks with up to six species and five antibiotics, we found many stable community topologies, all of which included

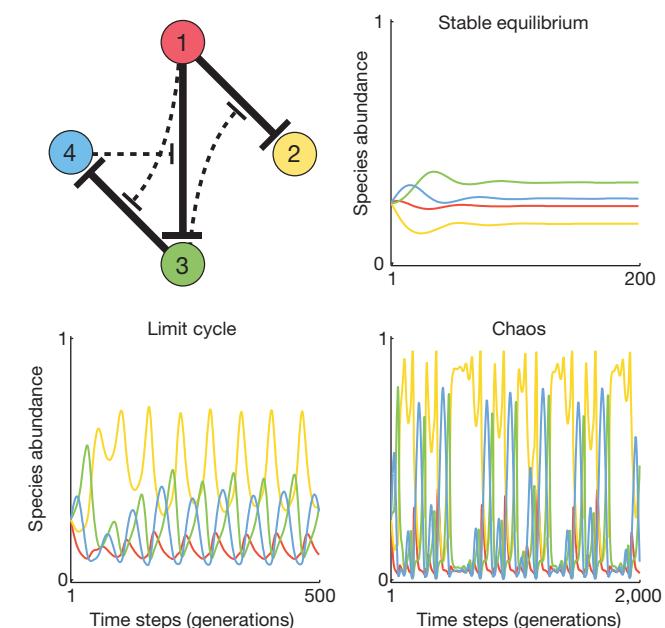


Figure 4 | Complex interaction networks support coexistence through diverse dynamic behaviours. An example of a stable network with four species interacting through production and degradation of three antibiotics. Depending on the strengths of antibiotic production and degradation, this community can coexist through stable equilibrium, limit cycles or chaos (Methods). Line colours indicate species identity.

antibiotic-degrading species (Extended Data Fig. 8). Notably, greater numbers of antibiotics generally increase the number of species that can coexist. While the fraction of network topologies that support stable communities can be small, the combinatorial increase in the number of possible networks provides multitudes of ways for antibiotic production and degradation to generate stable communities with large numbers of species. Coexistence among species in these various network topologies can occur through different dynamic behaviours, including stable fixed points, limit cycles, or chaos, depending on the strengths of antibiotic production K_P and degradation K_D (Fig. 4 and Extended Data Fig. 9). Thus, in addition to enabling the long-term coexistence of complex communities, the combined effect of antibiotic production and degradation can dramatically impact community dynamics.

Our findings suggest new possibilities for engineering multi-species microbial consortia³⁰ and shed light on the role of antibiotic production and degradation in maintaining biodiversity within natural microbial communities. Of course, natural microbial ecosystems are extremely complicated¹³; they contain orders of magnitude more species than we modelled and include additional interactions at the level of resource competition, metabolic cross-feeding, phage invasion, and predator-prey relationships. We expect that further insights into ecosystem stability and assembly will emerge by understanding how these mechanisms generate both pairwise and higher-order multi-species interactions.

Online Content Methods, along with any additional Extended Data display items and Source Data, are available in the online version of the paper; references unique to these sections appear only in the online paper.

Received 26 November 2014; accepted 20 April 2015.

Published online 20 May 2015.

1. Hutchinson, G. E. The paradox of the plankton. *Am. Nat.* **95**, 137–145 (1961).
2. May, R. M. Will a large complex system be stable? *Nature* **238**, 413–414 (1972).
3. Levin, S. A. Dispersion and population interactions. *Am. Nat.* **108**, 207–228 (1974).
4. Armstrong, R. A. & McGehee, R. Competitive exclusion. *Am. Nat.* **115**, 151–170 (1980).
5. Levin, S. A. Community equilibria and stability, and an extension of the competitive exclusion principle. *Am. Nat.* **104**, 413–423 (1970).
6. Stewart, F. M. & Levin, B. R. Partitioning of resources and the outcome of interspecific competition: a model and some general considerations. *Am. Nat.* **107**, 171–198 (1973).
7. Levin, B. R. Coexistence of two asexual strains on a single resource. *Science* **175**, 1272–1274 (1972).
8. Shores, N., Hegreness, M. & Kishony, R. Evolution exacerbates the paradox of the plankton. *Proc. Natl Acad. Sci. USA* **105**, 12365–12369 (2008).
9. Durrett, R. & Levin, S. Spatial aspects of interspecific competition. *Theor. Popul. Biol.* **53**, 30–43 (1998).
10. Czárán, T. L., Hoekstra, R. F. & Pagie, L. Chemical warfare between microbes promotes biodiversity. *Proc. Natl Acad. Sci. USA* **99**, 786–790 (2002).
11. Kerr, B., Riley, M. A., Feldman, M. W. & Bohannan, B. J. M. Local dispersal promotes biodiversity in a real-life game of rock-paper-scissors. *Nature* **418**, 171–174 (2002).
12. Reichenbach, T., Mobilia, M. & Frey, E. Mobility promotes and jeopardizes biodiversity in rock-paper-scissors games. *Nature* **448**, 1046–1049 (2007).
13. Hibbing, M. E., Fuqua, C., Parsek, M. R. & Peterson, S. B. Bacterial competition: surviving and thriving in the microbial jungle. *Nature Rev. Microbiol.* **8**, 15–25 (2010).
14. Chait, R., Vetsigian, K. & Kishony, R. What counters antibiotic resistance in nature? *Nature Chem. Biol.* **8**, 2–5 (2012).
15. Conlin, P. L., Chandler, J. R. & Kerr, B. Games of life and death: antibiotic resistance and production through the lens of evolutionary game theory. *Curr. Opin. Microbiol.* **21**, 35–44 (2014).
16. Green, J. & Bohannan, B. J. M. Spatial scaling of microbial biodiversity. *Trends Ecol. Evol.* **21**, 501–507 (2006).
17. Vetsigian, K., Jajoo, R. & Kishony, R. Structure and evolution of *Streptomyces* interaction networks in soil and in silico. *PLoS Biol.* **9**, e1001184 (2011).
18. Cordero, O. X. et al. Ecological populations of bacteria act as socially cohesive units of antibiotic production and resistance. *Science* **337**, 1228–1231 (2012).
19. Raynaud, X. & Nunan, N. Spatial ecology of bacteria at the microscale in soil. *PLoS ONE* **9**, e87217 (2014).
20. Perlin, M. H. et al. Protection of *Salmonella* by ampicillin-resistant *Escherichia coli* in the presence of otherwise lethal drug concentrations. *Proc. R. Soc. Lond. B* **276**, 3759–3768 (2009).
21. Wright, G. D. Bacterial resistance to antibiotics: enzymatic degradation and modification. *Adv. Drug Deliv. Rev.* **57**, 1451–1470 (2005).
22. Wootton, J. T. Indirect effects in complex ecosystems: recent progress and future challenges. *J. Sea Res.* **48**, 1–16 (2002).
23. Ohgushi, T., Schmitz, O. & Holt, R. D. *Trait-Mediated Indirect Interactions*. (Cambridge Univ. Press, 2012).
24. Bollenbach, T., Quan, S., Chait, R. & Kishony, R. Nonoptimal microbial response to antibiotics underlies suppressive drug interactions. *Cell* **139**, 707–718 (2009).
25. Dugatkin, L. A., Perlin, M., Lucas, J. S. & Atlas, R. Group-beneficial traits, frequency-dependent selection and genotypic diversity: an antibiotic resistance paradigm. *Proc. R. Soc. Lond. B* **272**, 79–83 (2005).
26. Yurtsev, E. A., Chao, H. X., Datta, M. S., Artemova, T. & Gore, J. Bacterial cheating drives the population dynamics of cooperative antibiotic resistance plasmids. *Mol. Syst. Biol.* **9**, 683 (2013).
27. Lenski, R. E. & Hattingh, S. E. Coexistence of two competitors on one resource and one inhibitor: a chemostat model based on bacteria and antibiotics. *J. Theor. Biol.* **122**, 83–93 (1986).
28. Hsu, S. B. & Waltman, P. A survey of mathematical models of competition with an inhibitor. *Math. Biosci.* **187**, 53–91 (2004).
29. Levin, S. A., Segel, L. A. & Adler, F. R. Diffuse coevolution in plant-herbivore communities. *Theor. Popul. Biol.* **37**, 171–191 (1990).
30. Großkopf, T. & Soyer, O. S. Synthetic microbial communities. *Curr. Opin. Microbiol.* **18**, 72–77 (2014).

Supplementary Information is available in the online version of the paper.

Acknowledgements We thank R. Ward, S. Levin, M. Elowitz, G. Bumim, A. Amir, S. Kryazhimskiy, Y. Gerardin, J. Meyer, N. Yin, E. Bairey, H. Chung and A. Palmer for critical feedback and discussions; M. Baym, M. Fischbach, M. Traxler, R. Chait, L. Stone and E. Gontag for strains. E.D.K. acknowledges government support under and awarded by the Department of Defense, Office of Naval Research, National Defense Science and Engineering Graduate (NDSEG) Fellowship, 32 CFR 168a. J.Z. acknowledges support from the Research Science Institute at MIT during the summer of 2011. K.V. acknowledges start-up funds from University of Wisconsin-Madison. R.K. acknowledges the support of a James S. McDonnell Foundation 21st Century Science Initiative in Studying Complex Systems Research Award 220020169, National Institutes of Health Grant R01GM081617, European Research Council Seventh Framework Programme ERC Grant 281891, Israeli Centers of Research Excellence I-CORE Program ISF Grant No. 152/11.

Author Contributions E.D.K., K.V. and R.K. designed the study; E.D.K., K.V. and J.Z. did the experiments; E.D.K. and K.V. did the simulations and modelling; E.D.K., K.V. and R.K. analysed the data and wrote the paper.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to K.V. (kalin@discovery.wisc.edu) and R.K. (rkishony@technion.ac.il).

METHODS

Mixed inhibition-zone model of cyclic three-species communities. *Model description.* We derived a formula for the changes of species abundances in the inhibition-zone model for the limit of complete mixing after each time step ($r_{\text{dispersal}} \rightarrow \infty$) and infinite population size. Let X_i be the abundance of species i per unit area, K_p the killing area around each producer and K_D the degradation area around each degrader. Since species disperse randomly, the fitness f_i of species i depends on its inherent growth rate g_i and the fraction of area $1 - P_{\text{kill}}$ in which it is not killed by antibiotics:

$$f_i = g_i(1 - P_{\text{kill}}).$$

For each antibiotic, a sensitive individual S will only be killed if it is located where there is at least one producer P within a K_p neighbourhood and there is no protecting antibiotic-degrading (R_D) species within a K_D neighbourhood. Since individuals are randomly placed, the number of cells within such neighbourhoods follows a Poisson distribution, so that if λ is the expected number of cells in a region then the probability of having zero cells in the region is $e^{-\lambda}$. Thus the probability of being in the killing zone of the producer is $1 - e^{-K_p X_p}$, and the probability of not being in the degradation zone of R_D is $e^{-K_D X_D}$, where X_p and X_D are respectively the abundances of species producing and degrading the antibiotic. Combining these we obtain

$$P_{\text{kill}} = e^{-K_D X_D} (1 - e^{-K_p X_p})$$

Thus for the cyclic three-species network the fitness of each species is:

$$f_1 = g_1 (1 - e^{-K_D X_2} (1 - e^{-K_p X_3}))$$

$$f_2 = g_2 (1 - e^{-K_D X_3} (1 - e^{-K_p X_1}))$$

$$f_3 = g_3 (1 - e^{-K_D X_1} (1 - e^{-K_p X_2}))$$

We update species abundance for each time step using the discrete time formula:

$$X_i(t+1) = \frac{X_i(t)f_i(t)}{\sum_j X_j(t)f_j(t)}$$

where the denominator ensures that total species abundance remains constant. *Linear stability analysis.* We found the fixed point numerically for each set of parameters and then found the eigenvalues of the Jacobian computed at the fixed point. The logarithm of the maximum absolute value of the eigenvalues was used as a stability measure in Fig. 2b, c to determine the region of stability. The fixed point is stable if this value is less than zero.

Robustness to mutation and invasion by cheating species. Each cheating species was cloned from a parent by copying all antibiotic phenotypes except for a mutation in the phenotype with respect to a single antibiotic: $P \rightarrow R_I$ for production cheaters, and $R_D \rightarrow S$ for degradation cheaters. Starting with a three-species community at equilibrium, we added a small amount of the cheating species ($X_{\text{cheater}} = 0.001$, with abundances of the other species renormalized so that $\sum X_i = 1$). We set the growth rate of the cheater to the parental growth rate with a growth advantage ϵ , so that $g_{\text{cheater}} = (1 + \epsilon)g_{\text{parent}}$. We ran each simulation until species abundances reached steady state or stable oscillations, and then we calculated the minimum and maximum values for each cheater during the last third of the simulation time.

Coexistence of communities with three species and two antibiotics. Cyclic three-species communities maintain coexistence even when the fastest growing species ceases antibiotic production, resulting in a simpler three-species, two-antibiotic community. This network exhibits robust coexistence similar to the three-antibiotic network, so long as the species that does not produce antibiotics is the fastest growing species (Extended Data Fig. 6a). This simpler community also resists invasion by any remaining production or degradation cheaters (Extended Data Fig. 7b).

We investigated whether further reductions of this network support coexistence. While coexistence is possible when only one of the two antibiotics is degraded, this coexistence occurs only for a small range of parameter values (Extended Data Fig. 6b). Based on a comprehensive sampling of parameter space, we found that further reductions to communities that produce only a single antibiotic do not support coexistence when inherent growth rates differ. This is consistent with the intuition that when only one antibiotic is being produced there will always be two species that are equivalent with respect to the antibiotic, which will lead to extinction of the slower-growing species. Thus, in contrast to spatial models where coexistence is possible on a single antibiotic¹¹, coexistence in well-mixed models requires three species and at least two antibiotics.

Generalizing the mixed inhibition-zone model for arbitrary interaction networks. *Model description.* We calculated the probability of each species being inhibited for a general ecosystem with n_s species and n_a antibiotics, where species

can be any of the four antibiotic phenotypes (P, S, R_I or R_D) for each antibiotic. The assumptions are the same as for the mixed inhibition-zone model of cyclic three-species communities. Consider the situation in which antibiotic production is stronger than degradation ($K_p > K_D$, the other scenario follows similarly). For an individual of a target sensitive species, each of the other $n_s - 1$ species can be either near enough to affect the target through production and degradation of some of the antibiotics (with probability P_{near}), can be at intermediate distances in which they affect the target through production but not degradation ($P_{\text{intermediate}}$), or can be far away so that they do not affect the target in any way (P_{far}). These three distance ranges make for $3^{n_s - 1}$ possible combinations for the presence or absence of the other species within any neighbourhood of the target species. Based on the phenotypes of the sensitive target species and the other species, we calculated all combinations that result in antibiotic inhibition. For each species i , we calculated the probabilities P_{near} , $P_{\text{intermediate}}$ and P_{far} from the levels of production and degradation of the antibiotics, in this case:

$$P_{\text{near}} = 1 - e^{-K_D X_i}$$

$$P_{\text{far}} = e^{-K_D X_i}$$

$$P_{\text{intermediate}} = P_{\text{far}} - P_{\text{near}}$$

Combining these, we calculated the probability P_c of each inhibitory combination occurring, multiplying the probabilities of each species being either near, intermediate, or far according to the requirements of each combination. Finally, we obtained the total probability of the target species being inhibited by summing across all possible inhibitory combinations with $P_{\text{kill}} = \sum_c P_c$. As in the three-species model we let the fitness of each species be:

$$f_i = g_i(1 - P_{\text{kill}})$$

and

$$X_i(t+1) = \frac{X_i(t)f_i(t)}{\sum_j X_j(t)f_j(t)}$$

Testing for coexistence in more complex networks. We randomly sampled up to 10^6 networks for each combination of 4–6 initial species and 1–5 initial antibiotics by randomly choosing one of the four phenotypes (S, P, R_I or R_D) for each species/antibiotic combination. When the total number of possible networks was less than 10^6 we tested all networks. We excluded networks that contained antibiotics or species with identical properties from subsequent simulations, since these are equivalent to instances of networks with a smaller number of antibiotics or species. With n_a initial antibiotics and n_s initial species at initial abundances $1/n_s$, we calculated the number of final species surviving at abundance $X_i > 0.01/n_s$ after 10,000 time steps. Simulations were binned by the maximum number of surviving species across all possible combinations of the parameters: $g_{\text{max}} = (1.2 \text{ or } 2.5)$, $K_p = (10 \text{ or } 30)$ and $K_D = (3 \text{ or } 10)$. For a given g_{max} species growth rates were evenly spaced between 1.0 and g_{max} . With total number of network topologies $n_{\text{total}} = 4^{n_a n_s} / n_a!$, even this sparse sampling of parameters shows that the total number of topologies that support stability is quite large (Extended Data Fig. 8).

Calculating effective number of species using Shannon diversity. We calculated effective species numbers by the Shannon entropy H of the species distribution:

$$H = - \sum_i X_i \log_2 X_i$$

with species frequencies X_i normalized such that $\sum_i X_i = 1$. Effective number of species is then given by 2^H .

Chemostat model of cyclic three-species communities. *Model description.* Our chemostat model is adapted from classical models of antibiotic interactions with a constant inflow of a single resource and constant dilution²⁸. Ordinary differential equations track the concentration of the resource (Z), the abundances of three species (X_i), and the concentrations of three antibiotics (C_i), implementing the cyclical interaction network as described previously. Each species consumes resources and produce antibiotics. The species degrade antibiotics to which they are resistant, and their growth is inhibited by antibiotics to which they are sensitive. The resource dynamic is:

$$\frac{dZ}{dt} = (Z_o - Z)D - \sum_i \frac{X_i G_i}{m},$$

where Z_o is the concentration of the added resource, D is the dilution rate, G_i a function that describes the growth rate of species i considering antibiotic and resource levels, and m is a conversion factor between resources and species,

so that G_i/m is the resource consumption rate of each species. The species abundances change according to:

$$\frac{dx_i}{dt} = X_i(G_i - D) \text{ with } G_i = g_i \frac{Z}{k_z + Z} e^{-K_p C_j}.$$

Growth rate increases with the concentration of resources following Monod kinetics, where k_z determines the concentration of half-maximal resource absorption. The growth rate is also affected by antibiotics, decreasing exponentially with the level of antibiotics to which a species is sensitive (here an antibiotic with concentration C_j inhibits species i). The parameter K_p characterizes the strength of antibiotic inhibition, and g_i is a species-specific maximal growth rate.

The antibiotic concentration dynamics are given by:

$$\frac{dC_j}{dt} = pX_jG_j - K_D X_k C_j - DC_j,$$

The parameter p determines the amount of antibiotics produced per cell division, the second term assumes that species k degrades the antibiotic produced by species j through mass action kinetics, and the parameter K_D characterizes the strength of antibiotic degradation.

Rescaled chemostat model. To analyze the stability of the three-species three-antibiotic interaction network, we first note that

$$\frac{d}{dt}(\sum_i \frac{X_i}{m} - Z_o + Z) = -D(\sum_i \frac{X_i}{m} - Z_o + Z),$$

and therefore after a transient

$$Z = Z_o - \sum_i X_i/m$$

which allows us to eliminate the resource concentration degree of freedom.

We rescale variables to reduce the number of free parameters: time so that $D = 1$, Z and k_z so that $Z_o = 1$, and we eliminate m by substituting $X_i' = X_i/m$, $C_i' = C_i/m$, $K_p' = mK_p$ and $K_D' = mK_D$. Dropping the primes we obtain:

$$\begin{aligned} \frac{dx_i}{dt} &= X_i(G_i - 1) \\ G_i &= g_i \frac{1 - \sum_j X_j}{k_z + 1 - \sum_j X_j} e^{-K_p C_j} \\ \frac{dC_j}{dt} &= pX_jG_j - K_D X_k C_j - C_j \end{aligned}$$

Linear stability analysis. The problem of finding the fixed point in the rescaled model was reduced to a single nonlinear algebraic equation, which was numerically solved for different parameters to determine the corresponding equilibria for X_i and C_j . At the fixed point $\frac{dx_i}{dt} = 0$, and therefore $G_i = 1$. This allows us to express the C_j variables as functions of

$$\Theta = \frac{1 - \sum_i X_i}{k_z + 1 - \sum_i X_i}$$

and the g_i constants. Substituting the expressions for C_j into

$$pX_jG_j - K_D X_k C_j - C_j = 0$$

results in a system of linear equations for X_i terms, which can be solved as a function of Θ . This allows us to express $\sum_i X_i$ as a function of Θ and thus obtain a nonlinear equation for Θ . For Extended Data Fig. 5a, the stability of the fixed point was determined by finding the eigenvalues of the Jacobian at the fixed point. A negative largest real component of the eigenvalues indicates stability.

Parameter values for main text figures: Fig. 1: grids were 200×200 with wrapping boundaries. Species were initially randomly seeded on the grid at low total abundance (1%), or in an initial configuration with three large domains for $r_{\text{dispersal}} > 20$ to reduce the initial chance of extinction of all species at high levels of mixing (with intrinsic resistance). Fig. 1e, lines show average of ten simulations, calculating the effective species numbers based on Shannon diversity of species within 40×40 sub-regions (shown in insets) at 100 evenly spaced locations on the grid at the final time; other parameters: $r_{\text{production}} = 3$, $r_{\text{degradation}} = 3$, $r_{\text{dispersal}} = (3, 4, 6, 10, 20, 80, 200)$; runtimes for each $r_{\text{dispersal}}$ are respectively: (150, 150, 75, 75, 75, 25, 25), after which the overall spatial patterns were relatively unchanging. For intrinsic resistance with $r_{\text{dispersal}} = (80, 200)$ it was possible for all species to be inhibited at the same time step, in which case we chose one cell at random to proceed to the next generation.

Fig. 2: we estimated the size of the basin of attraction by running simulations starting from 210 different initial species abundances that were equally spaced on a triangle lattice inside the $x_1 + x_2 + x_3 = 1$ two-dimensional simplex. The basin of attraction area is the fraction of simulations that move to the fixed point with all three species coexisting. Fig. 2a, $g_1 = 1.0$, $g_2 = 1.1$, $g_3 = 1.2$; left: $K_p = 4$, $K_D = 0$;

right: $K_p = 12$, $K_D = 3$. Fig. 2b, left: $K_p = 4$, $K_D = 0$; right: $K_p = 40$, $K_D = 4$. Fig. 2c, $g_1 = 1$, $g_2 = 3$, $g_3 = 2$.

Fig. 3: $g_1 = 1.0$, $g_2 = 1.1$, $g_3 = 1.2$, $K_p = 16$, $K_D = 4$.

Fig. 4: $g_1 = 1.0$, $g_2 = 1 + 1/30$, $g_3 = 1 + 2/30$, $g_4 = 1.1$; stable equilibrium: $K_p = 10$, $K_D = 10$; limit cycles: $K_p = 12$, $K_D = 20$; chaos: $K_p = 20$, $K_D = 55$.

Code availability. MATLAB code for the spatial inhibition-zone model, mixed inhibition-zone model and chemostat model is available from the authors upon request.

Note on the mechanism of three-way attenuating interactions. Antibiotic degradation is a common mechanism of antibiotic resistance²¹ and is therefore a probable mechanism for the attenuation of inhibition that we observed in the three-species interaction assays and the antibiotic modulation assays (Extended Data Figs 1 and 2). There is strong evidence in support of this assumption based on many reports of antibiotic degrading species protecting sensitive species from antibiotics^{20,25,26,31–33}. Antibiotic degradation is also a known mechanism of resistance for the antibiotics in which we observed significant attenuation (Extended Data Fig. 1c): tobramycin³⁴, ciprofloxacin³⁵, β -lactam antibiotics³⁶ (cefoxitin, piperacillin, penicillin), nitrofurantoin³⁷, rifampin³⁸ and trimethoprim³². Furthermore, we experimentally confirmed that degradation through β -lactamases contributes to the attenuation of β -lactam antibiotics through a modified version of the antibiotic modulation assays that uses β -lactamase inhibitors (Extended Data Figs 1d and 2c). However, we note that the models we use can be easily generalized to additional mechanisms that generate antibiotic attenuation, for example through antibiotic suppression^{24,39}.

Interaction experiments. No statistical methods were used to predetermine sample size. *Three-species interaction assays.* All possible pairings of 16 *Streptomyces* species from a spore collection (~ 100 spores per μl) were pinned 1 cm apart on 9-cm Petri dishes with $1 \times$ oatmeal agar, and plates were grown in a darkened plastic bin for 13 days at room temperature. Overlay media without thiamine was prepared in test tubes in 5-ml aliquots, melted in an autoclave (to melt, not sterilize), and cooled to 42 °C in a water bath. An overnight liquid-broth culture of *E. coli* expressing YFP was mixed 9:1 with 1,000× thiamine stock, 50 μl of the *E. coli* mixture was added to overlay agar, quickly vortexed to mix and then poured uniformly onto the plate. Plates were incubated at 37 °C for 18 h and photographed in bright field and YFP channels to record the antibiotic inhibition zone. Combinations in which one species created a large inhibition zone were scored using the modulation index $M = (r_m - r_o)/(r_m + r_o)$, where r_m is the radius of the inhibition zone in the direction of the modulator and r_o is the radius of the inhibition zone in the opposite direction. Extended Data Fig. 1b shows data from 54 combinations in which one species made a large inhibition zone, with five antibiotic producers and 11 modulator species (one plate removed due to contamination). Extended Data Fig. 2a shows images and scoring. We observed similar attenuation in follow-up technical replicas of individual strain combinations and when testing on other media.

Antibiotic modulation assays. Antibiotic modulation assays were similar to the three-species assays except that the antibiotic producing species was replaced with pure antibiotic and the assay geometry was modified to enable more high-throughput measurements. Modulator inoculation plates were prepared by spreading 50 μl of 20% glycerol spore stocks onto $1/2 \times$ oatmeal agar 9-cm Petri dishes with glass beads, incubated for 1 week at 30 °C in plastic bins, then stored at room temperature. 25 ml of $1/2 \times$ oatmeal agar was spread evenly over a 14-cm Petri dish. On day 1, a small amount of antibiotic stock was pipetted onto the centre of the agar, absorbed and stored at 30 °C (cefoxitin 10 μl , chloramphenicol 10 μl , ciprofloxacin 10 μl , doxycyclin 10 μl , nitrofurantoin 20 μl , penicillin 20 μl , piperacillin 15 μl , rifampicin 20 μl , tobramycin 40 μl , trimethoprim 20 μl). On day 2, a thin radial line for each of three modulator species was inoculated using the edges of sterilized 60-mm cover glass (VWR 48393-070) onto the 14-cm Petri dish. Plates were incubated in plastic bins at 30 °C for 3 days and small areas of contamination were cut out of the agar. Overlay media was prepared in 25-ml aliquots, as in the three-species assays. Overnight cultures of yellow fluorescent protein (YFP) and cyan fluorescent protein (CFP) producing *E. coli* were combined at equal ratios (by OD) and mixed 9:1 with 1,000× thiamine stock. 1 ml of the *E. coli* mixture was added to the overlay media, vortexed, and poured evenly over the surface of the oatmeal agar. All plates were incubated at 37 °C between 16–21 h and photographed in bright field, YFP and CFP channels to record the antibiotic inhibition zone. Test strains included the two strongest attenuators from the three-species assay, *Streptomyces coelicolor* and a collection of soil isolates from Massachusetts and Colorado soils. Combinations were marked with a dot for significance in Extended Data Fig. 1c where the average modulation index M deviated from 0 by at least twice the standard error in the mean ($P < 0.05$, two-sided *t*-test, $n = 3$ technical replicas). These images were also visually inspected to confirm warping of the inhibition-zone boundary. Extended Data Fig. 2b shows examples and scoring for different antibiotics.

β -Lactamase inhibitor assays. 1 ml of 1 mM β -lactamase inhibitor was spread evenly over 14-cm 1/2 \times oatmeal plates with glass beads and allowed to dry in a ventilation hood (to minimize contamination). Overlay with YFP and CFP *E. coli* was the same as the antibiotic modulation assay except that 1 ml of 1 mM β -lactamase inhibitor stock was added to the overlay media and vortexed before adding *E. coli*. Extended Data Fig. 1d shows data for a collection of 14 species, four antibiotics and three β -lactamase inhibitors. Extended Data Fig. 2c shows example assay images.

Isolation of soil bacteria. Isolation and identification of five *Streptomyces* from soil A (Massachusetts) as described in ref. 17. We chose five strains from a collection of 47 isolates that did not produce strong antibiotic inhibition zones on oatmeal agar. Soil B (Colorado) was collected from White River National Forest in unincorporated Pitkin County. We chose ten strains from a collection of 55 isolates that did not produce strong inhibition zones, including two *Streptomyces* species, five Actinobacteria species and three non-Actinobacteria species. Species identification was through Sanger sequencing of 16S ribosomal DNA.

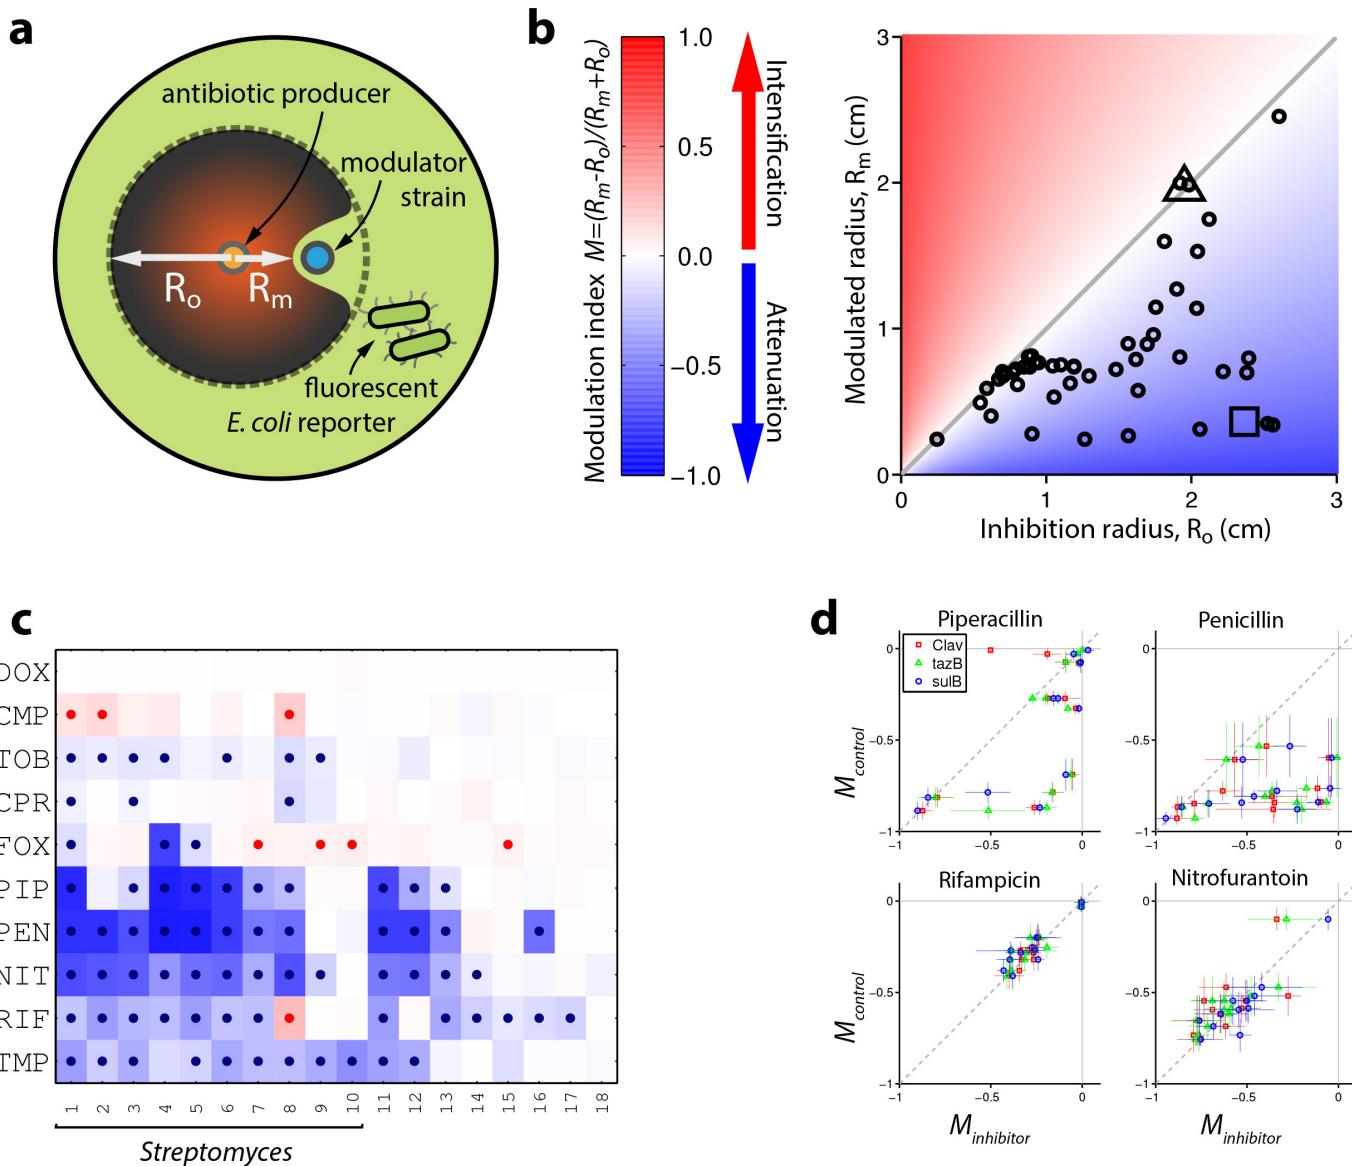
Imaging and data analysis. The three-species assay images were photographed using a custom-built automated fluorescent imaging system as described in ref. 40. Other assays were using a similar but newer system. Levels of antibiotic modulation were scored by manually recording the edges of the inhibition zones of the fluorescent *E. coli* using MATLAB. Example images were normalized to fill the dynamic range and gamma adjusted to improve contrast.

Growth media. 1 \times oatmeal agar: 72.5 g l⁻¹ Difco oatmeal agar (BD). 1/2 \times oatmeal agar: 31.25 g l⁻¹ Difco oatmeal agar (BD) and 6.25 g l⁻¹ Bacto agar (BD), with ~100 μ l l⁻¹ 10 M NaOH to adjust to pH 7. Media for overlays consisted of M63 salts (2 g l⁻¹ (NH₄)₂SO₄, 13.6 g l⁻¹ KH₂PO₄, 0.5 mg l⁻¹ FeSO₄·7H₂O) supplemented with 0.4% glucose, 0.02% casamino acids, 1 mM MgSO₄, 1.5 mM thiamine, 7.5 g l⁻¹ bacto-agar (BD), adjusting pH to 7.0 with 1 M NaOH. 1,000 \times thiamine was prepared at 1.5 M in H₂O and filter sterilized.

Chemical stocks. Cefoxitin (Sigma C-4786) 50 mg ml⁻¹ in H₂O, chloramphenicol (Sigma C0378) 30 mg ml⁻¹ in EtOH, ciprofloxacin (Sigma 17850) 10 mg ml⁻¹ in

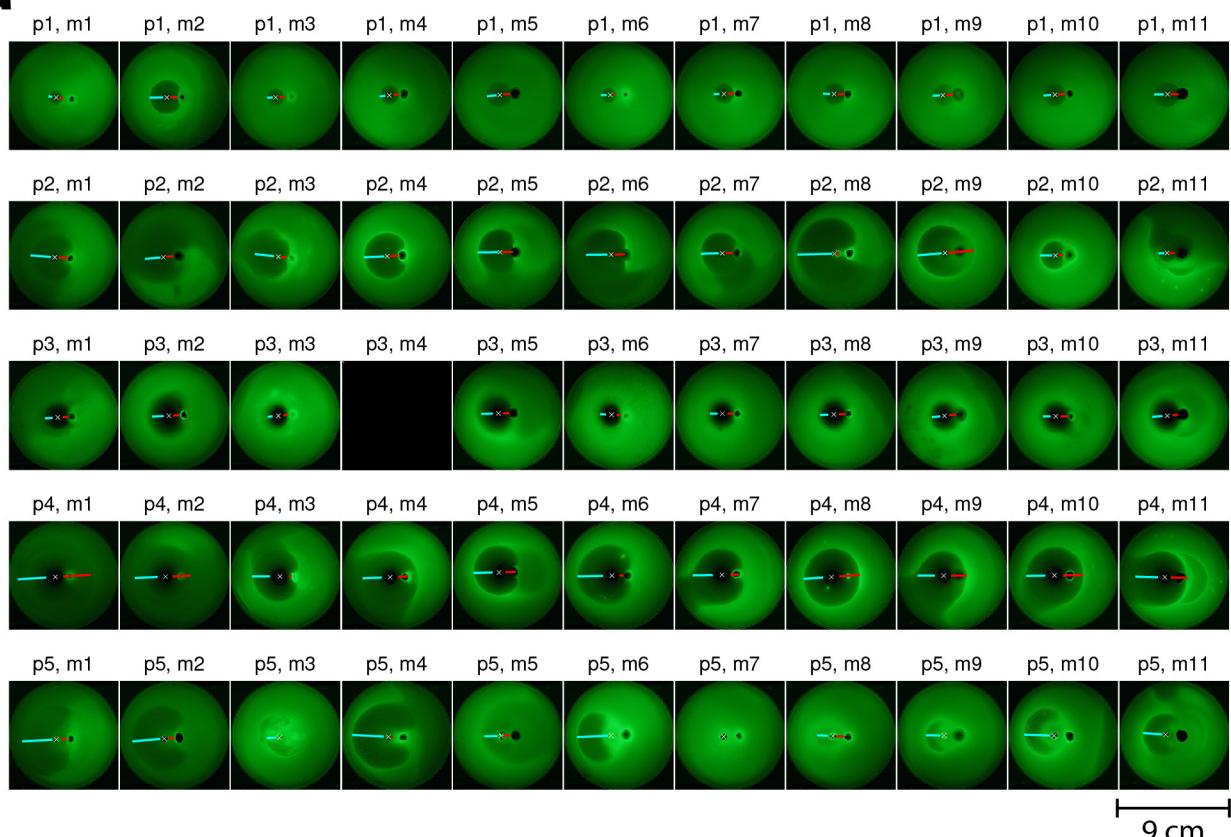
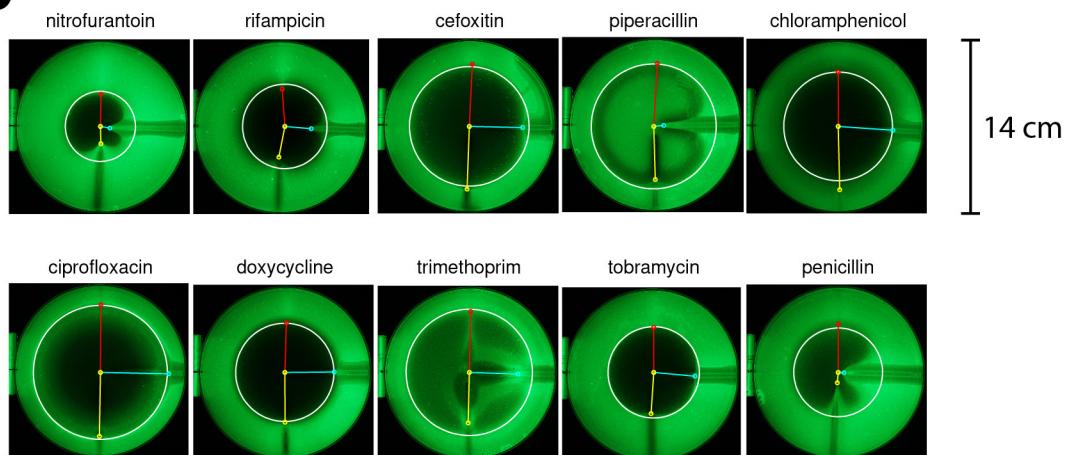
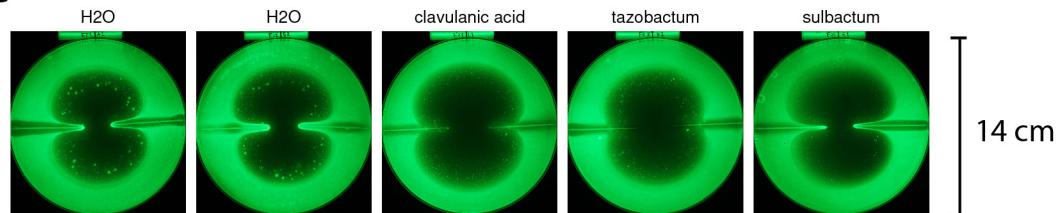
H₂O + 3 μ l ml⁻¹ 10 M HCl, doxycyclin (Sigma D-9891) 50 mg ml⁻¹ in H₂O, nitrofurantoin (Sigma N7878) 10 mg ml⁻¹ in DMF, penicillin (Sigma 13750) 50 mg ml⁻¹ in H₂O, piperacillin (Sigma P8396) 50 mg ml⁻¹ in H₂O, rifampicin (Sigma R3501) 16 mg ml⁻¹ in DMSO, tobramycin (Sigma T4014) 50 mg ml⁻¹ in H₂O, and trimethoprim (Sigma T7883) 5 mg ml⁻¹ in H₂O. β -Lactamase inhibitors were prepared at 1 mM in H₂O: clavulanic acid (as potassium clavulanate, Sigma 33454), tazobactam sodium salt (Sigma T2820), sulbactam (Sigma S9701).

31. D'Costa, V. M. Sampling the antibiotic resistome. *Science* **311**, 374–377 (2006).
32. Walsh, F. & Duffy, B. The culturable soil antibiotic resistome: a community of multi-drug resistant bacteria. *PLoS ONE* **8**, e65567 (2013).
33. Brook, I. The role of beta-lactamase-producing-bacteria in mixed infections. *BMC Infect Dis.* **9**, 202 (2009).
34. Shaw, K. J., Rather, P. N., Hare, R. S. & Miller, G. H. Molecular genetics of aminoglycoside resistance genes and familial relationships of the aminoglycoside-modifying enzymes. *Microbiol. Rev.* **57**, 138–163 (1993).
35. Wetzstein, H. G., Stadler, M., Tichy, H. V., Dalhoff, A. & Karl, W. Degradation of ciprofloxacin by basidiomycetes and identification of metabolites generated by the brown rot fungus *Gloeophyllum striatum*. *Appl. Environ. Microbiol.* **65**, 1556–1563 (1999).
36. Livermore, D. M. β -Lactamases in laboratory and clinical resistance. *Clin. Microbiol. Rev.* **8**, 557–584 (1995).
37. Rafii, F. & Hansen, E. B. Isolation of nitrofurantoin-resistant mutants of nitroreductase-producing *Clostridium* sp. strains from the human intestinal tract. *Antimicrob. Agents Chemother.* **42**, 1121–1126 (1998).
38. Hoshino, Y. et al. Monoxygenation of rifampicin catalyzed by the rox gene product of *Nocardia farcinica*: structure elucidation, gene identification and role in drug resistance. *J. Antibiot. (Tokyo)* **63**, 23–28 (2010).
39. Chait, R., Craney, A. & Kishony, R. Antibiotic interactions that select against resistance. *Nature* **446**, 668–671 (2007).
40. Chait, R., Shrestha, S., Shah, A. K., Michel, J.-B. & Kishony, R. A differential drug screen for compounds that select against antibiotic resistance. *PLoS ONE* **5**, e15179 (2010).



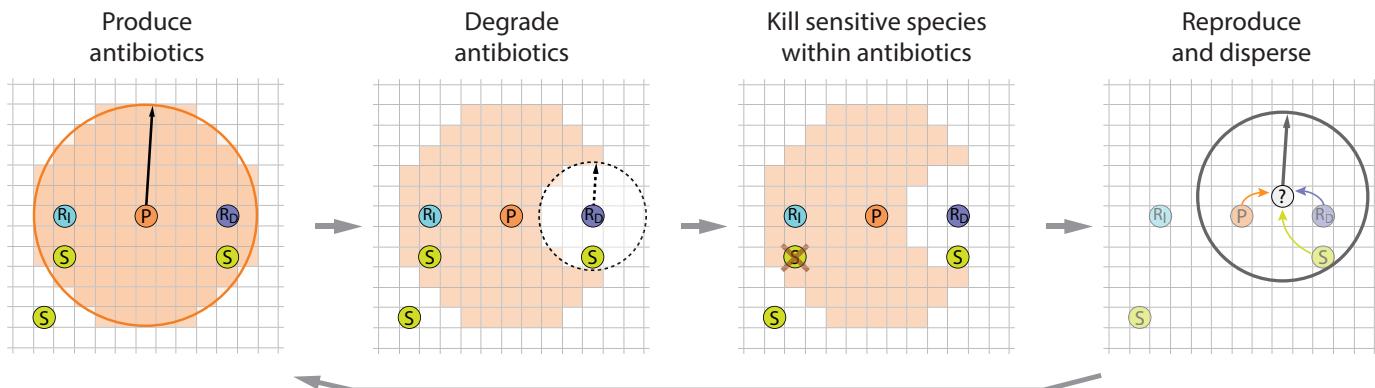
Extended Data Figure 1 | Antibiotic attenuation is widespread among natural antibiotic producing species isolated from soil. **a**, Diagram of three-species assays: measuring the antibiotic inhibition zone of *E. coli* around a producer strain enables quantification of three-species interactions caused by a modulator strain. **b**, Attenuating interactions dominate among a set of 54 *Streptomyces* producer–modulator combinations. Triangle and square markers show scoring from example images in Fig. 1b, left and right respectively. **c**, Antibiotic modulation assays: attenuating interactions dominate among combinations of soil species with a panel of ten pure antibiotics. Combinations are coloured by modulation index as in **b** and marked with a dot where the modulation index is significantly different from zero ($n=3$, Methods). Strains 1 and 2 are the strongest modulators from the three-species assays, strain three is

Streptomyces coelicolor, strains 4–18 are additional soil isolates (Extended Data Table 1). DOX, doxycycline; CMP, chloramphenicol; TOB, tobramycin; CPR, ciprofloxacin; FOX, cefoxitin; PIP, piperacillin; PEN, penicillin; NIT, nitrofurantoin; RIF, rifampicin; TMP, trimethoprim. **d**, Scatter plots of average modulation index M for 12 species from panel c, with and without β -lactamase inhibitors. Points occur off the diagonal for the β -lactams piperacillin and penicillin but near the diagonal for the structurally unrelated antibiotics rifampicin and nitrofurantoin, consistent with attenuation of the β -lactams through a mechanism of antibiotic degradation. Error bars are standard error of the mean for technical replicas with inhibitor ($n=3$) or without inhibitor (control, $n=6$). Control is addition of H_2O instead of inhibitor. Clav, clavulanic acid; tazB, tazobactum; sulB, sulbactam.

a**b****c**

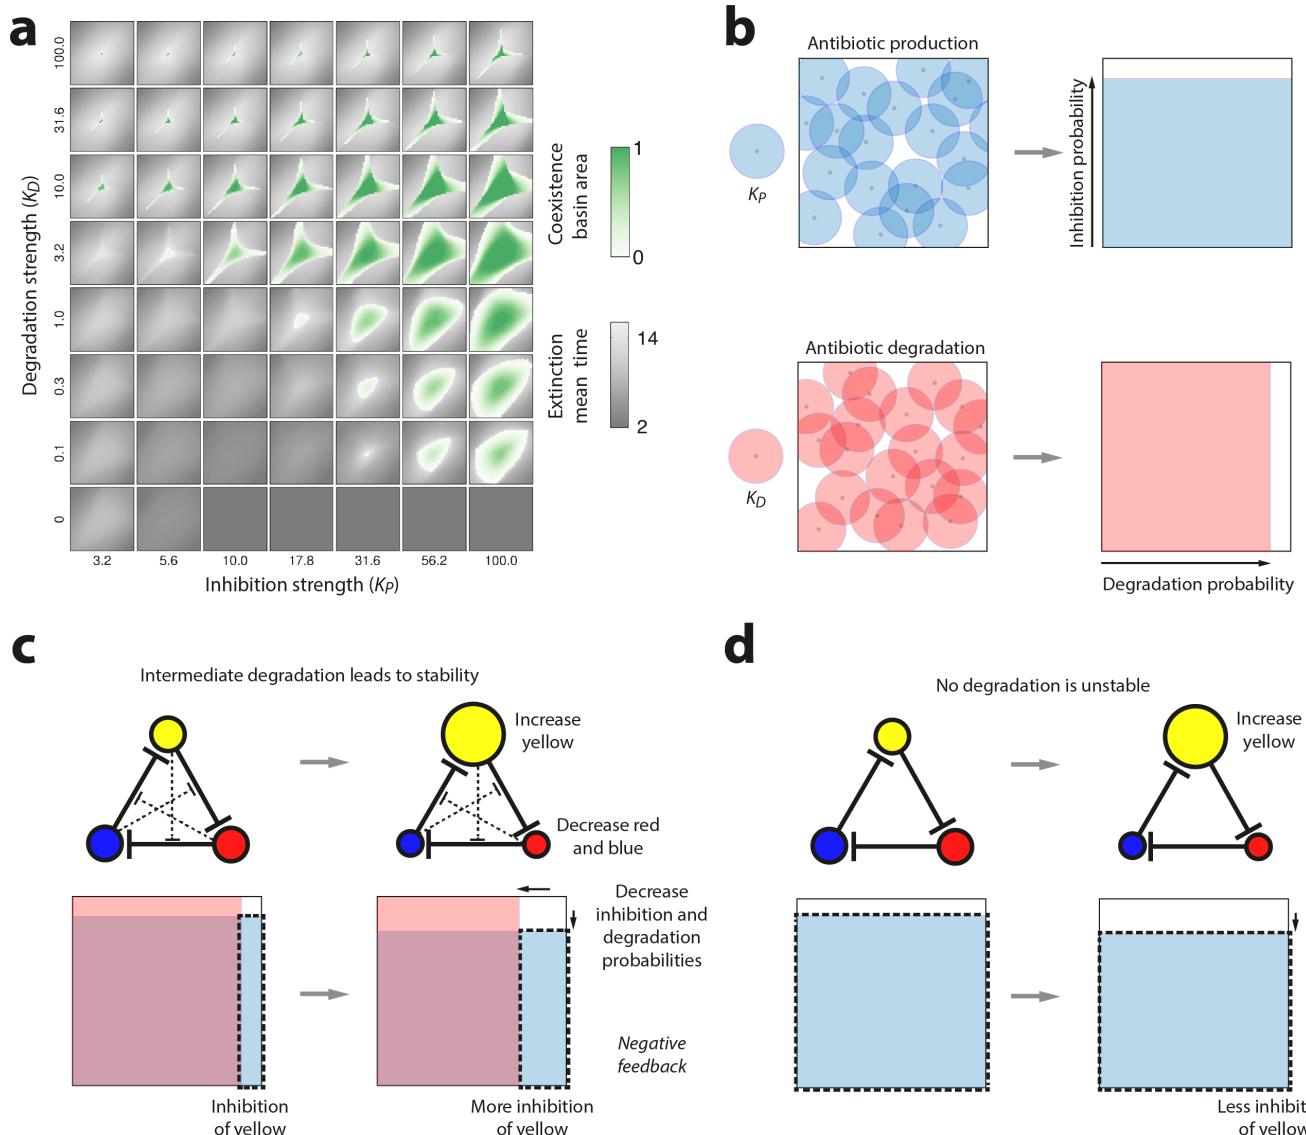
Extended Data Figure 2 | Example images from three-species interaction assays and antibiotic assays. **a**, Images and scoring for three-species interaction assays from Extended Data Fig. 1b. p indicates producer species, m indicates potential modulator species. **b**, Example images and scoring from the antibiotic modulation assays from Extended Data Fig. 1c. Each plate shows tests for modulation of antibiotic inhibition for three different species against a different antibiotic. Colour lines show the size of the inhibition zone at the

location of the radially positioned modulator species; the white circles show the radius of the zone of inhibition as inferred from the left side of the plate that contains no modulators. **c**, Example images from an antibiotic modulation assay with β -lactamase inhibitors and the β -lactam antibiotic cefoxitin. Left and right side of each plate is inoculated with a line of species four from Extended Data Fig. 1c. Attenuation is significantly reduced by the β -lactamase inhibitors (especially clavulanic acid and tazobactam) when compared to controls (H_2O).


Extended Data Figure 3 | Illustration of the spatial inhibition-zone model.

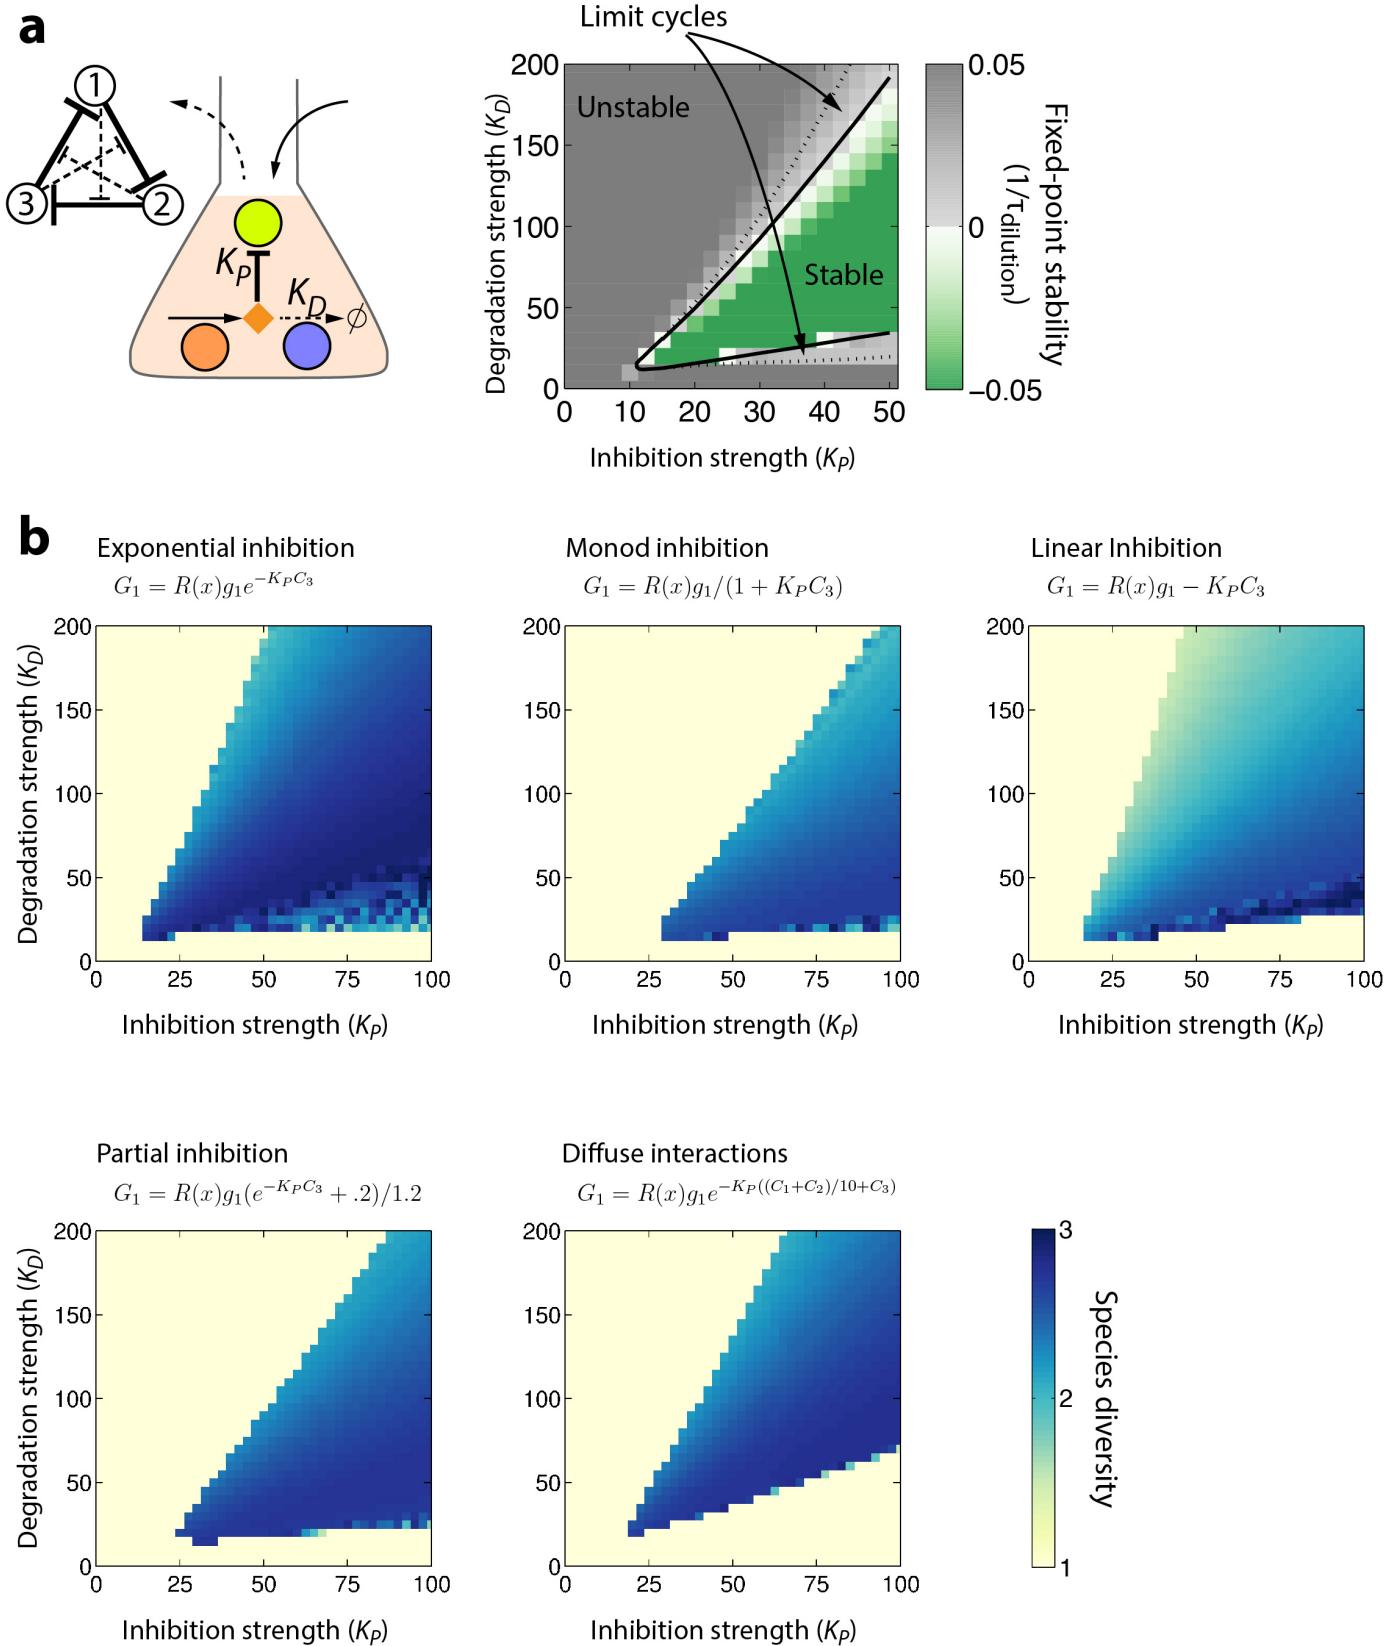
The simulation is performed on a grid of size $L \times L$. A single individual occupies each grid location. During each generation: (1) Individuals from species of type P produce antibiotics within a circle of radius $r_{\text{production}}$; (2) individuals of type R_D remove antibiotics within a circle of radius $r_{\text{degradation}}$; (3) all sensitive

individuals (type S) are killed at any locations that still contain the corresponding antibiotics (crossed out S; antibiotic values at each location are calculated at centre positions); (4) empty locations of a new grid are filled by randomly choosing from any surviving individuals within a radius $r_{\text{dispersal}}$. If there are no surviving individuals within $r_{\text{dispersal}}$ then an empty location remains empty.



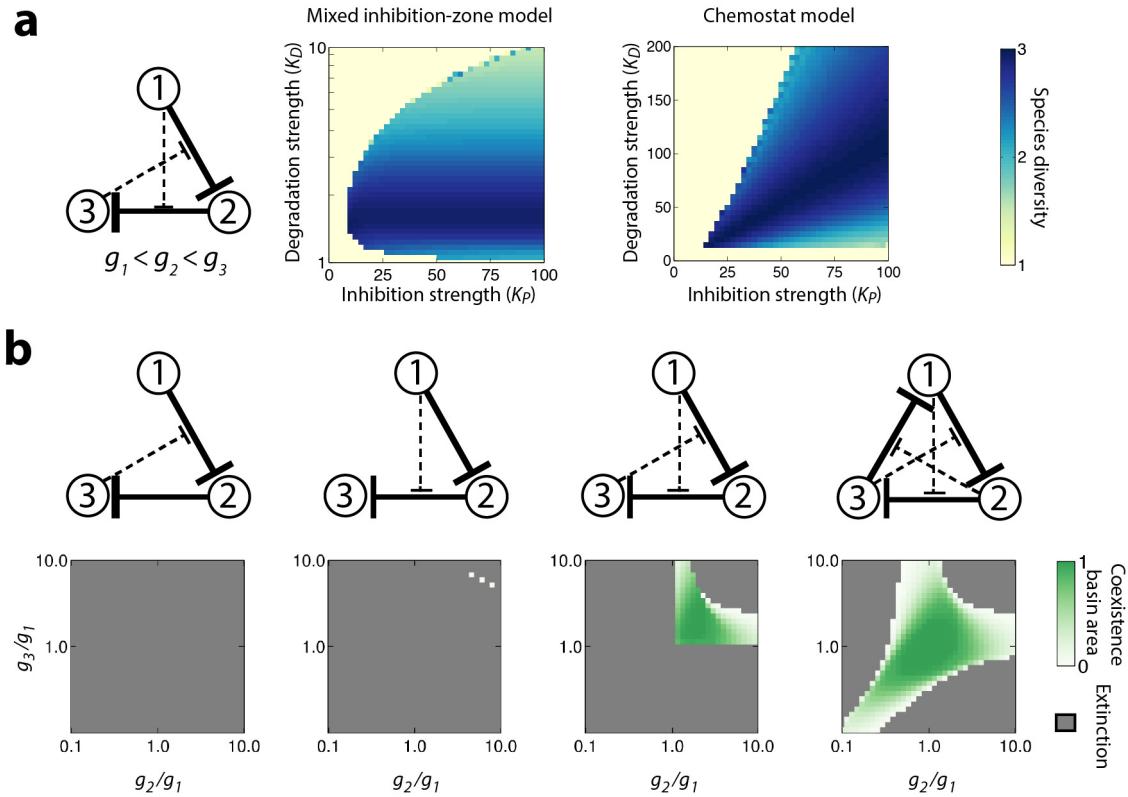
Extended Data Figure 4 | Dependence of coexistence on degradation in the mixed inhibition-zone model. **a**, Stability analysis of the full parameter space in the mixed inhibition-zone model. Using simulations, we tested the stability of cyclic three-species, three-antibiotic communities with dense sampling of all possible parameters for the inhibition-zone model, varying strengths of K_P and K_D , initial abundances of species 1–3 and growth ratios g_2/g_1 and g_3/g_1 . As in Fig. 2b, each grid shows a 100-fold range of growth rate ratios from 0.1 to 10. Large basins of attraction exist across a wide range of parameter values, with maximal stability at high levels of antibiotic production and intermediate levels of degradation. **b–d**, Intuition for why coexistence depends on degradation. **b**, The inhibition-zone model calculates the probability of a given sensitive

species being inhibited by an antibiotic producer (blue), or being protected by a degrading species (red), given the relative strengths of production (K_P) and degradation (K_D). The expected area covered by the overlapping circles (left) is used to calculate the corresponding inhibition and attenuation probabilities (percentage area of filled boxes, right; Methods). **c**, Focusing on one antibiotic in a stable three-species community, increasing the abundance of the yellow species creates negative feedback by decreasing the abundance of blue and red, which results in more inhibition of yellow. **d**, Communities are not stable at low levels of degradation due to positive feedback, whereby increasing the abundance of the yellow species results in a decrease of inhibition.



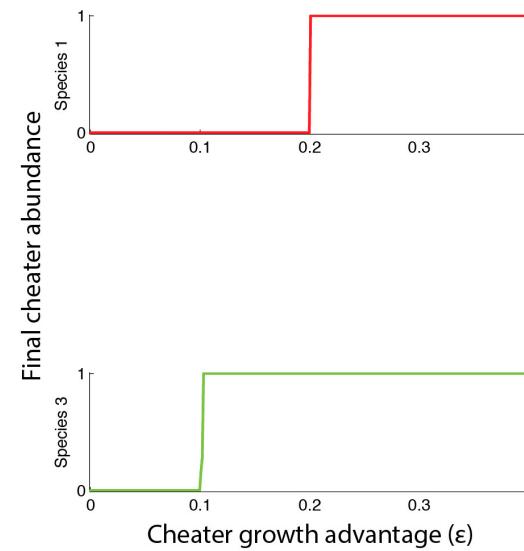
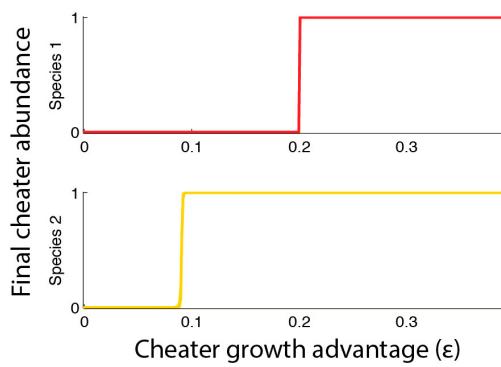
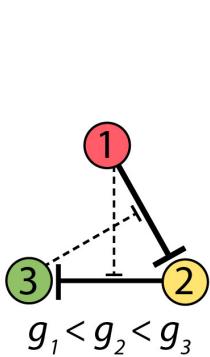
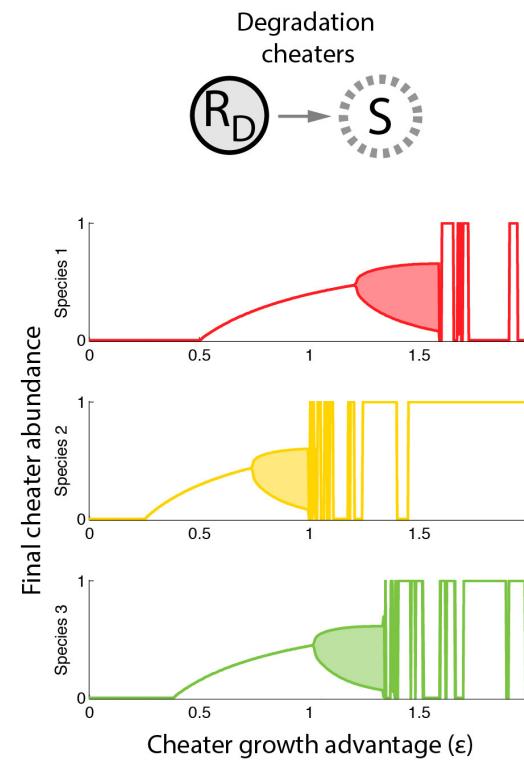
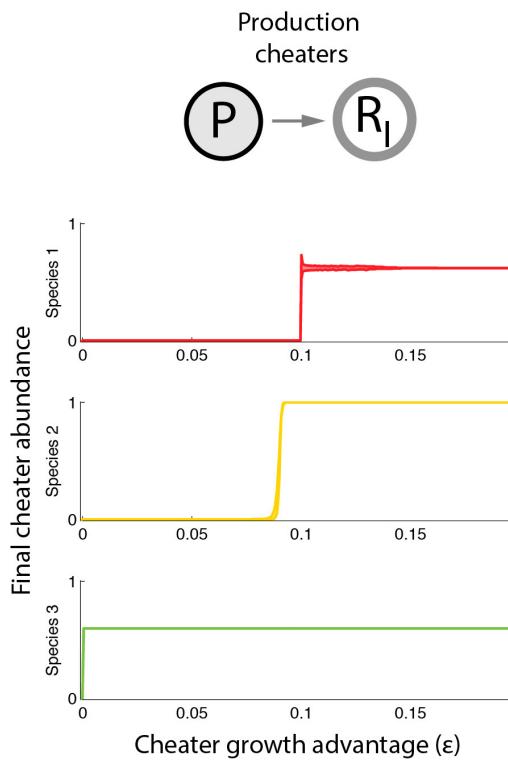
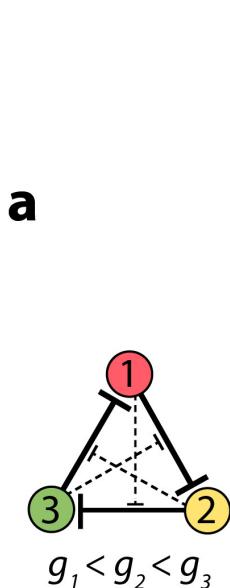
Extended Data Figure 5 | Coexistence of antibiotic degrading communities in a well-mixed chemostat model with three antibiotics. **a**, Stability of the three-species, three-antibiotic community in a single resource chemostat model in which species and antibiotics are completely homogeneous. Communities in the chemostat model may coexist through fixed points or limit cycles. For parameter sets in which the fixed point of the chemostat model was unstable, we started simulations close to the fixed point to determine if the community coexists through a limit cycle. Limit cycles occur in the areas between the dashed and solid black lines. Chemostat parameters: $g_1 = 3, g_2 = 6, g_3 = 9, k_z = 0.5, P = 1$. **b**, Communities coexist in the chemostat for a wide range of assumptions regarding antibiotic mechanisms. We simulated the chemostat model while changing how the action of the antibiotics is modelled and

observed robust coexistence across all models. For each simulation we started all species at equal concentration ($X_i = 0.2$), ran the simulation until $t = 100$ generations and calculated the Shannon diversity of the final species levels. The default chemostat model assumes exponential inhibition of species by antibiotics, but similar coexistence is observed for Monod-like inhibition, for linear inhibition (for $G_i < 0$ we set $G_i = 0$), when species are only partially inhibited or when each species is sensitive at some level to all antibiotics. G_1 is the growth rate of species 1 under inhibition, while g_1 is its maximal rate of growth; $R(x)$ captures resource dependence on current species levels; C_1, C_2 and C_3 are the concentrations of antibiotics produced by species 1–3 respectively (Methods). Equations for the growth of species 2 and 3 have the same form as G_1 . All other parameters are the same as for panel **a**.



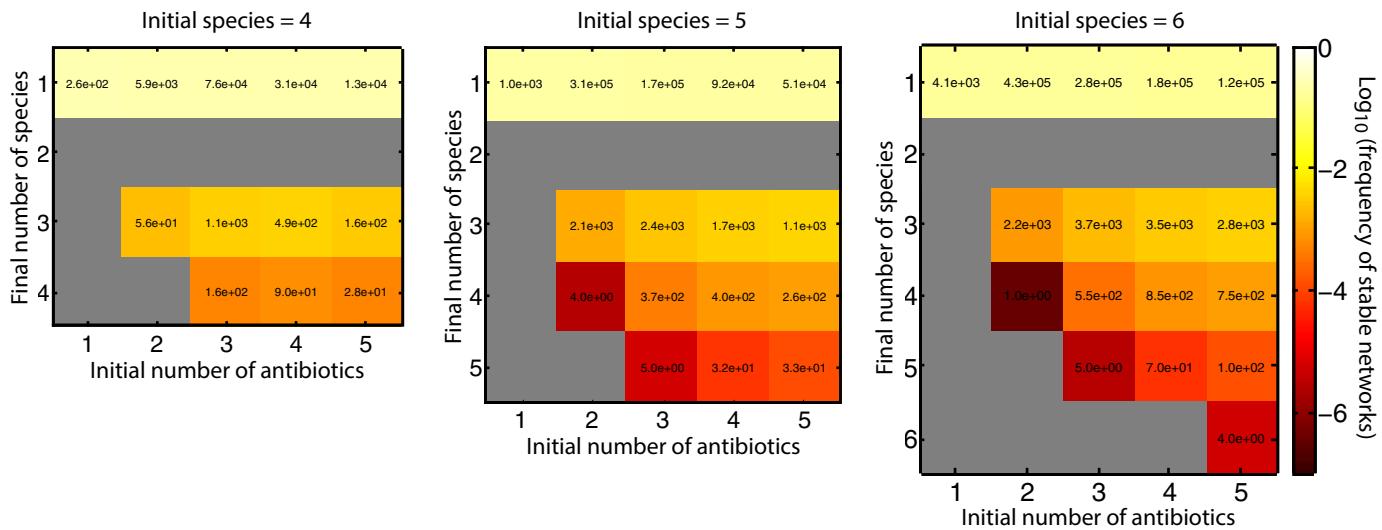
Extended Data Figure 6 | Coexistence of communities with three species and two antibiotics. **a**, Comparing community diversity in the mixed inhibition-zone model and the chemostat model. For each simulation we started all species at equal concentration ($X_i = 1/3$ for the inhibition-zone model, $X_i = 0.2$ for the chemostat model), ran the simulation for until $t = 100$ generations and calculated the Shannon diversity of the final species levels. Other parameters are the same as in Fig. 2 for the inhibition-zone model or Extended Data Fig. 5 for the chemostat model. **b**, Three species communities

require two antibiotics for stability. When only one antibiotic is degraded the community either lacks stability (first panel), or is stable only for a small number of growth rates and initial conditions (second panel). When two antibiotics are degraded the community is robustly stable to differences in species growth rates and initial conditions (third panel), provided that the antibiotics inhibit the faster-growing species (species 2 and 3). Basin colours as in Fig. 2b; grey shows parameters for which no initial conditions were stable; $K_P = 40$, $K_D = 4$.



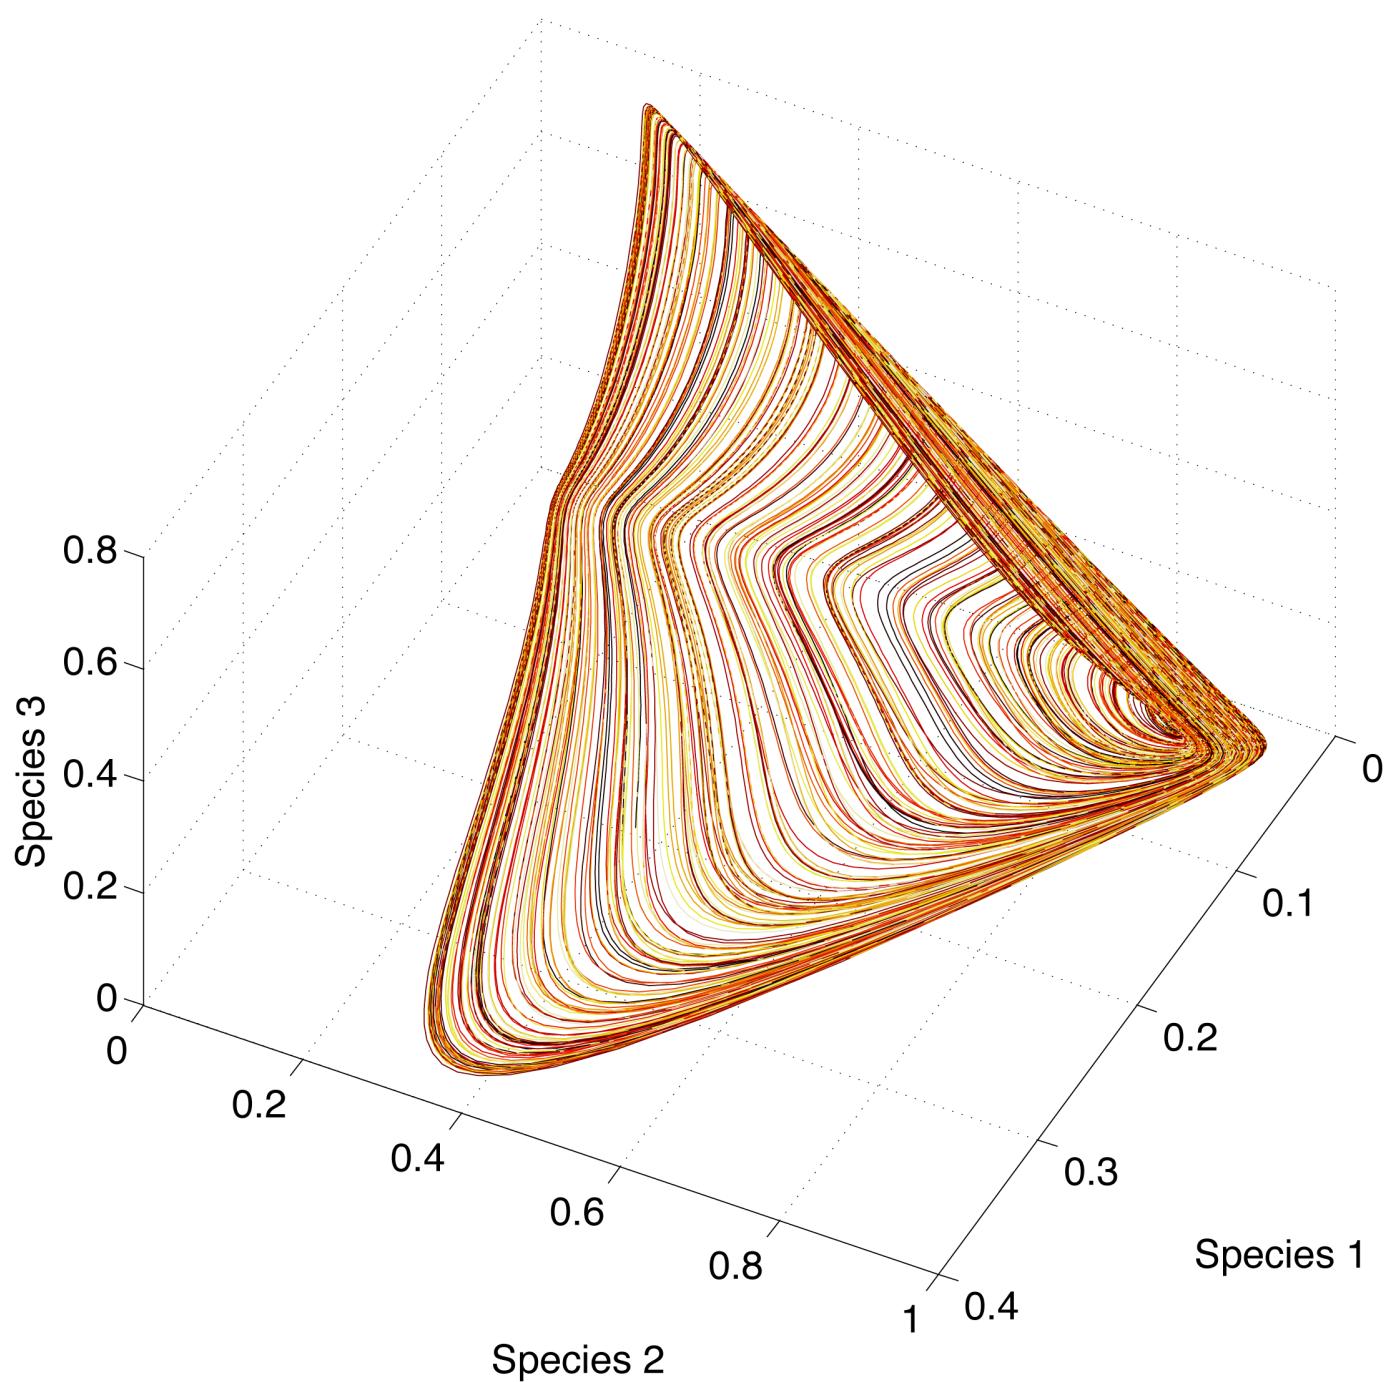
Extended Data Figure 7 | Robustness of three species communities to invasion by cheaters in the mixed inhibition-zone model. Analysis of production cheaters ($P \rightarrow R_I$, left) and degradation cheaters ($R_D \rightarrow S$, right). As in Fig. 3, we plot the final abundance of each cheater as a function of its growth advantage ε over its parent species. **a**, Cheaters cannot invade the three-species, three-antibiotic network when their growth advantage is small, except for the production cheater of the species with the fastest inherent growth rate (species 3, green line), which replaces its parent generating a new stable community of three species interacting through two antibiotics. **b**, This resulting three-species, two-antibiotic community is resilient to invasion by all cheaters; cheaters must have a substantial growth advantage to invade

and take over. Parameters for both networks are the same as Fig. 3. Shaded areas indicate the maximum and minimum abundance when the community reaches stable oscillations. Note, the analysis above is for networks with $g_1 < g_2 < g_3$. The alternative network of $g_1 > g_2 > g_3$, is less robust to invasion by production cheaters. Two cheaters can invade this community even with small ϵ : the production cheater for species 2 invades the network in **a** and gives rise to a stable three-species two-antibiotic community, while the production cheater for species 1 can take over the communities shown in **a** and **b**. For degradation cheaters this alternative network is similarly robust to cheaters as shown.



Extended Data Figure 8 | Complex network topologies support coexistence of larger numbers of species in the mixed inhibition-zone model. For given initial numbers of species and antibiotics, sets of up to 10^6 communities with random networks were simulated and the final number of surviving species recorded. The number inside each square shows how many networks resulted in the specified number of final species (after removing networks that did not

use all of the initial antibiotics, Methods). Colours show the frequency of each outcome within all simulated networks, with grey where no stable networks were found. We sparsely sampled parameters for species growth rates and antibiotic production and degradation levels (Methods). The sparse sampling means that a given network topology may exhibit stability for parameter combinations that were not tested.



Extended Data Figure 9 | A community with chaotic dynamics. Plotting the abundance of species 1–3 for the network from Fig. 4b. We show the last 30,000

steps of a simulation with 40,000 total generations, coloured with a slowly changing gradient. The trajectories form a strange attractor.

Extended Data Table 1 | Strain information for experimental assays

ED Fig. 1b ED Fig. 2	ED Fig. 1c	Species	Notes	Assays
p1		<i>Kutzneria sp.</i> 744	producer	1
p2		<i>Streptomyces clavuligerus</i> (ATCC 27064)	producer	1
p3		<i>Streptomyces sviceus</i> (ATCC 29083)	producer	1
p4		<i>Streptomyces sp.</i> Mg1	producer	1
p5		<i>Streptomyces albus</i> (J1074)	producer	1
m1		<i>Streptomyces SPB74</i>		1
m2		<i>Streptomyces griseoflavus</i> Tu4000		1
m3		<i>Streptomyces sp.</i> AA#4		1
m4	1	<i>Streptomyces pristinispiralis</i> (ATCC 25486)		1, 2, 3
m5		<i>Streptomyces hygroscopicus</i> (ATCC 53653)		1
m6		<i>Streptomyces SPB78</i>		1
m7		<i>Streptomyces ghanaensis</i> (ATCC 14672)		1
m8	2	<i>Streptomyces roseosporus</i> (NRRL 11379)		1, 2, 3
m9		<i>Streptomyces viridochromogenes</i> (DSM 40736)		1
m10		<i>Streptomyces lividans</i> (TK24)		1
m11		<i>Streptomyces sp.</i> E14		1
	3	<i>Streptomyces coelicolor</i>		2, 3
	4	<i>Streptomyces sp.</i> G1-4, JN020492	soil A	2, 3
	5	<i>Kitasatospora sp.</i> G4-12, JN020549	soil A	2, 3
	6	<i>Streptomyces sp.</i> G2-6, JN020526	soil A	2, 3
	7	Unclassified <i>Streptomycetaceae</i>	soil B	2, 3
	8	<i>Streptomyces sp.</i> G4-8, JN020545	soil A	2, 3
	9	<i>Streptomyces sp.</i> G1-14	soil A	2
	10	<i>Streptomyces sp.</i> RI18-DHHV3	soil B	2
	11	<i>Amycolatopsis sp.</i>	soil B, Actinobacteria	2, 3
	12	<i>Kribbella sp.</i> NBRC 104208	soil B, Actinobacteria	2, 3
	13	<i>Rhodococcus sp.</i>	soil B, Actinobacteria	2, 3
	14	<i>Microbacterium sp.</i>	soil B, Actinobacteria	2
	15	<i>Rhodococcus sp.</i> 602	soil B, Actinobacteria	2
	16	<i>Agromonas sp.</i>	soil B, Alphaproteobacteria	2, 3
	17	<i>Lysobacter sp.</i>	soil B, Gammaproteobacteria	2
	18	<i>Bradyrhizobiaceae</i> Bosea	soil B, Alphaproteobacteria	2
S		<i>E. coli</i> MC4100-YFP/PCS-I	assay strain w/ chromosomal YFP	1
S		<i>E. coli</i> MC4100, pZS2R-YFP	assay strain w/ YFP plasmid	2, 3
S		<i>E. coli</i> MC4100, pZS2R-CFP	assay strain w/ CFP plasmid	2, 3

Assay numbers: 1, three-species interaction assay; 2, antibiotic modulation assay; 3, β -lactamase inhibitor assay.

Global genetic analysis in mice unveils central role for cilia in congenital heart disease

You Li^{1*}, Nikolai T. Klena^{1*}, George C. Gabriel^{1*}, Xiaoqin Liu^{1*}, Andrew J. Kim¹, Kristi Lemke¹, Yu Chen¹, Bishwanath Chatterjee¹, William Devine², Rama Rao Damerla¹, Chienfu Chang¹, Hisato Yagi¹, Jovenal T. San Agustin³, Mohamed Thahir^{4,5}, Shane Anderton¹, Caroline Lawhead¹, Anita Vescovi¹, Herbert Pratt⁶, Judy Morgan⁶, Leslie Haynes⁶, Cynthia L. Smith⁶, Janan T. Eppig⁶, Laura Reinholdt⁶, Richard Francis¹, Linda Leatherbury⁷, Madhavi K. Ganapathiraju^{4,5}, Kimimasa Tobita¹, Gregory J. Pazour³ & Cecilia W. Lo¹

Congenital heart disease (CHD) is the most prevalent birth defect, affecting nearly 1% of live births¹; the incidence of CHD is up to ten-fold higher in human fetuses^{2,3}. A genetic contribution is strongly suggested by the association of CHD with chromosome abnormalities and high recurrence risk⁴. Here we report findings from a recessive forward genetic screen in fetal mice, showing that cilia and cilia-transduced cell signalling have important roles in the pathogenesis of CHD. The cilium is an evolutionarily conserved organelle projecting from the cell surface with essential roles in diverse cellular processes. Using echocardiography, we ultrasound scanned 87,355 chemically mutagenized C57BL/6J fetal mice and recovered 218 CHD mouse models. Whole-exome sequencing identified 91 recessive CHD mutations in 61 genes. This included 34 cilia-related genes, 16 genes involved in cilia-transduced cell signalling, and 10 genes regulating vesicular trafficking, a pathway important for ciliogenesis and cell signalling. Surprisingly, many CHD genes encoded interacting proteins, suggesting that an interactome protein network may provide a larger genomic context for CHD pathogenesis. These findings provide novel insights into the potential Mendelian genetic contribution to CHD in the fetal population, a segment of the human population not well studied. We note that the pathways identified show overlap with CHD candidate genes recovered in CHD patients⁵, suggesting that they may have relevance to the more complex genetics of CHD overall. These CHD mouse models and >8,000 incidental mutations have been sperm archived, creating a rich public resource for human disease modelling.

Inbred C57BL/6J mice mutagenized with ethynitrosourea (ENU) were bred to recover recessive coding mutations (Extended Data Fig. 1). Phenotyping was conducted using non-invasive fetal echocardiography, an ultrasound imaging modality also employed clinically for CHD diagnosis. This unbiased study design allows detection of even those CHD mutants inviable to term⁶, a segment of CHD cases that would be largely clinically inaccessible for genetic analysis. Using a combination of colour flow, two-dimensional, spectral Doppler and M-mode imaging, a wide spectrum of CHD phenotypes were recovered, including cardiac looping defects, cardiac outflow tract malalignment/septation defects, ventricular/atrioventricular septal defects, and various right/left heart obstructive lesions⁶ (Fig. 1 and Supplementary Videos 1, 2). The CHD diagnoses were confirmed with micro-computed tomography (CT)/micro-magnetic resonance imaging (MRI), necropsy and histopathology (Extended Data Fig. 1). Heritability was ascertained with the recovery of two or more G3 fetuses with the same CHD phenotype ($n = 4.6 \pm 2.6$). Ultrasound scanning of 87,355 G3 fetuses from 2,651 mutant lines (defined by the G1 sire) recovered 218 CHD mutant lines. These were curated in the Mouse Genome Informatics Database (<http://www.informatics.jax.org/searchtool/Search.do?query=b2b&submit=Quick%0D%0ASearch>), with G1 sperm cryopreserved in the Jackson Laboratory (JAXMice).

Many mutant lines (30%) exhibited complex CHD associated with heterotaxy—the randomization of left–right patterning. These heterotaxy mutants often died prenatally. We note that clinical studies show that 16% of human fetuses with complex CHD have heterotaxy⁷, and that these also exhibit high prenatal/intrauterine death (30–60%)^{8,9}. Also commonly observed in mutant lines with or without laterality defects are double outlet right ventricle (DORV) and atrioventricular septal defects (AVSD), phenotypes that usually resulted in prenatal lethality. We note that DORV and AVSD are also observed with high prevalence in aborted/stillborn human fetuses³. Most heterotaxy mutant lines exhibited three distinct visceral organ situs phenotypes: normal situs solitus, mirror symmetric situs inversus, or randomization with heterotaxy, with complex CHD observed only with heterotaxy (Extended Data Fig. 2). These three phenotypes are found in patients and mice with mutations causing primary ciliary dyskinesia (PCD)^{10,11}, a sinopulmonary disease arising from immotile/dyskinetic respiratory cilia. This reflects the common requirement for motile cilia in airway clearance and left–right patterning^{11,12}. Hence, we conducted videomicroscopy of the tracheal epithelia to assess cilia motility as a proxy for the node.

Mutation recovery was conducted with whole-exome sequencing analysis of a single mutant from each of 113 mutant lines. Cumulatively this yielded 7,996 coding mutations in 4,809 genes (Extended Data Fig. 3a). Genotyping analysis identified the pathogenic mutation as one consistently homozygous across all G3 mutants bearing the same phenotype. In mutant lines with multiple homozygous mutations on the same chromosome linked to the CHD phenotype, the pathogenic mutation was identified with segregation analysis of additional mutants, and/or based on phenotype comparison to existing knockout/mutant alleles. For five lines (b2b002, b2b012, b2b284, b2b1519 and b2b1146), the mutation was confirmed with mapping analyses¹³.

In this manner, 91 pathogenic mutations were recovered in 61 genes (Supplementary Data 1). The distribution of missense, splicing and nonsense mutations was similar to other ENU mutagenesis studies¹⁴ (Extended Data Fig. 3a). All mutations except one were in highly conserved amino acids (Supplementary Data 2). Five of the 19 splicing mutations are >2 bases from the canonical splice donor-acceptor sites, the farthest being 10 bases away (*Dnah5*^{c.13329-10T>A}; b2b1537; validated by complementary DNA sequencing, Extended Data Fig. 4). Multiple alleles were recovered in 15 genes (Extended Data Table 1). Of 61 CHD genes recovered, 27 were not previously annotated to cause CHD (Extended Data Fig. 3b), either because no mutant/knockout mice were available, or because the CHD phenotype was not investigated (Supplementary

¹Department of Developmental Biology, University of Pittsburgh School of Medicine, Pittsburgh, Pennsylvania 15201, USA. ²Department of Pathology, University of Pittsburgh School of Medicine, Pittsburgh, Pennsylvania 15261, USA. ³Program in Molecular Medicine, University of Massachusetts Medical School, Worcester, Massachusetts 01605, USA. ⁴Department of Biomedical Informatics, University of Pittsburgh School of Medicine, Pittsburgh, Pennsylvania 15206, USA. ⁵Intelligent Systems Program, School of Arts and Sciences, University of Pittsburgh, Pittsburgh, Pennsylvania 16260, USA. ⁶The Jackson Laboratory, Bar Harbor, Maine 04609, USA. ⁷The Heart Center, Children's National Medical Center, Washington DC 20010, USA.

*These authors contributed equally to this work.

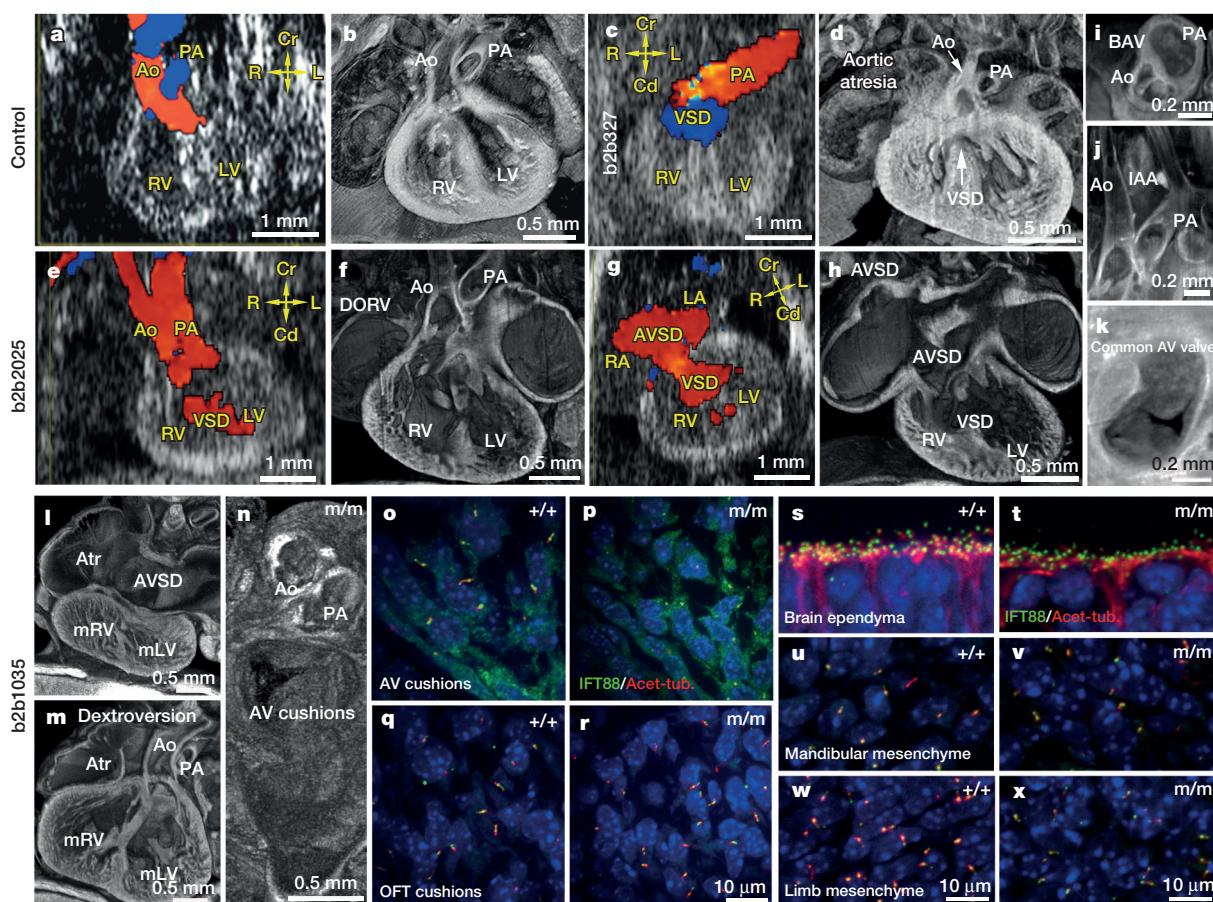


Figure 1 | Ultrasound diagnoses of CHD and cilia defects in CHD mutants. **a, b**, Vevo 2100 colour flow imaging showed criss-crossing of blood flow indicating normal aorta (Ao) and pulmonary artery (PA) alignment (**a**; Supplementary Video 1), confirmed by histopathology (**b**). Cd, caudal; Cr, cranial; L, left; LV, left ventricle; R, right; RV, right ventricle. **c, d**, Embryonic day (E)16.5 mutant mouse (line b2b327) exhibited a blood flow pattern indicating single great artery (pulmonary artery) and ventricular septal defect (VSD) (**c**), suggesting aortic atresia with ventricular septal defect, confirmed by histopathology (**d**). **e–h**, Colour flow imaging of E15.5 mutant mouse (line b2b2025) with heterotaxy (stomach on right; Supplementary Video 4) showed side by side aorta and pulmonary artery, with the aorta emerging from the right ventricle, indicating DORV/ventricular septal defect (**e, f**; Supplementary

Video 2) and the presence of AVSD (**g, h**; Supplementary Videos 3,4). LA, left atrium; RA, right atrium. **i–k**, Histopathology also showed a bicuspid aortic valve (BAV) (**i**), interrupted aortic arch (IAA) (**j**), and common atrioventricular (AV) valve (**k**). **l–n**, *Cc2d2a*-mutant mouse exhibits dextrocardia with ventricular inversion (dextroversion) (**m**), and AVSD (**l**) with malformed atrioventricular cushions (**n**), but normal outflow cushions. Atr, atrium; mLV, morphologic left ventricle; m/m, *Cc2d2a*-mutant mouse; mRV, morphologic right ventricle. **o–x**, Confocal imaging of E12.5 *Cc2d2a*-mutant mouse (m/m) versus wild-type (+/+) embryo sections showed no cilia in the atrioventricular cushion (**o, p**), but normal ciliation in the outflow cushion (OFT cushion) (**q, r**). **s–x**, Fewer and shorter cilia were observed in other mutant embryo tissues. Red, acetylated tubulin (Acet-tub.); green, IFT88.

Data 1). In some instances, the existing mutants died early, precluding CHD analysis, such as the *Smarca4*-knockout mouse with peri-implantation lethality¹⁵. A *Hectd1* mutant was previously described that only had neural tube defects¹⁶ versus our *Hectd1*^{b2b327} mutant, which was shown to have aortic atresia (Fig. 1c, d). The identical *Adamts6* mutation was found in three lines, and was shown to be a spontaneous mutation in the C57BL/6J support colony (Supplementary Data 1).

Thirty-four of the 61 CHD genes are cilia-related (Fig. 2), 23 in laterality-mutant lines (Extended Data Table 2 and Extended Data Figs 3a, 5). Twelve are required for motile cilia function, with eight clinically known to cause PCD and three that are probably novel PCD genes (*Ccdc151*, *Daw1* and *Pkd1l1*) (Extended Data Tables 2, 3 and Supplementary Data 1). A mutation was recovered in *Foxj1*, a transcription factor known to regulate ciliogenesis. While *Foxj1*-knockout mice have no cilia in the airway¹⁷, our *Foxj1*^{b2b774} mutant harbouring a missense mutation in the DNA-binding domain has ciliated trachea with robust ciliary motion (Supplementary Video 5), but with no nodal flow (Supplementary Video 6).

Mutations in 22 ciliome genes affect primary cilia—11 in laterality- and 11 in non-laterality-mutant lines (Extended Data Table 2 and Extended Data Figs 3a, 5). Among non-laterality mutants are *Cep110* (also

known as *Cntrl*), a centrosome component, *Jbts17*, a cilia transition zone protein, and four other genes (*Fuz*, *Myh10*, *Dctn5*, *Lrp2*) that also have important roles in vesicular/endocytic trafficking, a biological process essential for ciliogenesis¹⁸ (Fig. 2). Many genes affecting laterality are associated with non-PCD ciliopathies (*Anks6*, *Bicc1*, *Tbc1d32*, *Cep290*, *Dync2h1*, *Ift74*, *Ift140*, *Nek8*, *Pcsk5*, *Pskh1*, *Cc2d2a*, *Tmem67*; Extended Data Table 3)¹⁹. Analysis of a *Cc2d2a*-mutant embryo by confocal imaging showed shorter and fewer cilia throughout the embryo (Fig. 1s–x). In the *Cc2d2a*-mutant embryo, whereas cilia were abundant in the outflow cushions, they were absent in the atrioventricular cushions (Fig. 1p, r). Consistent with this, *Cc2d2a*-mutant embryos die prenatally with AVSD (Fig. 1l, n), but with normal semilunar valves (Fig. 1m, n).

Of the 26 non-ciliome CHD genes, 15 are known to mediate cell signalling, including the SHH signalling pathway (*Fuz*, *Kif7*, *Lrp2*, *Sufu*, *Tbc1d32*), WNT/planar cell polarity signalling (*Ptk7*, *Prickle1*, *Fuz*), TGF-β/BMP signalling (*Cfc1*, *Ltpb1*, *Megf8*, *Pcsk5*, *Smad6*, *Tab1*) and calcium signalling (*Pkd1l1*, *Pkd1*) (Fig. 2). Five of these were recovered in the laterality-mutant lines (*Tbc1d32*, *Cfc1*, *Megf8*, *Pcsk5*, *Pkd1l1*). These pathways are all transduced or modulated by the cilium^{20,21}, and have important roles in cardiovascular development and/or left-right patterning.

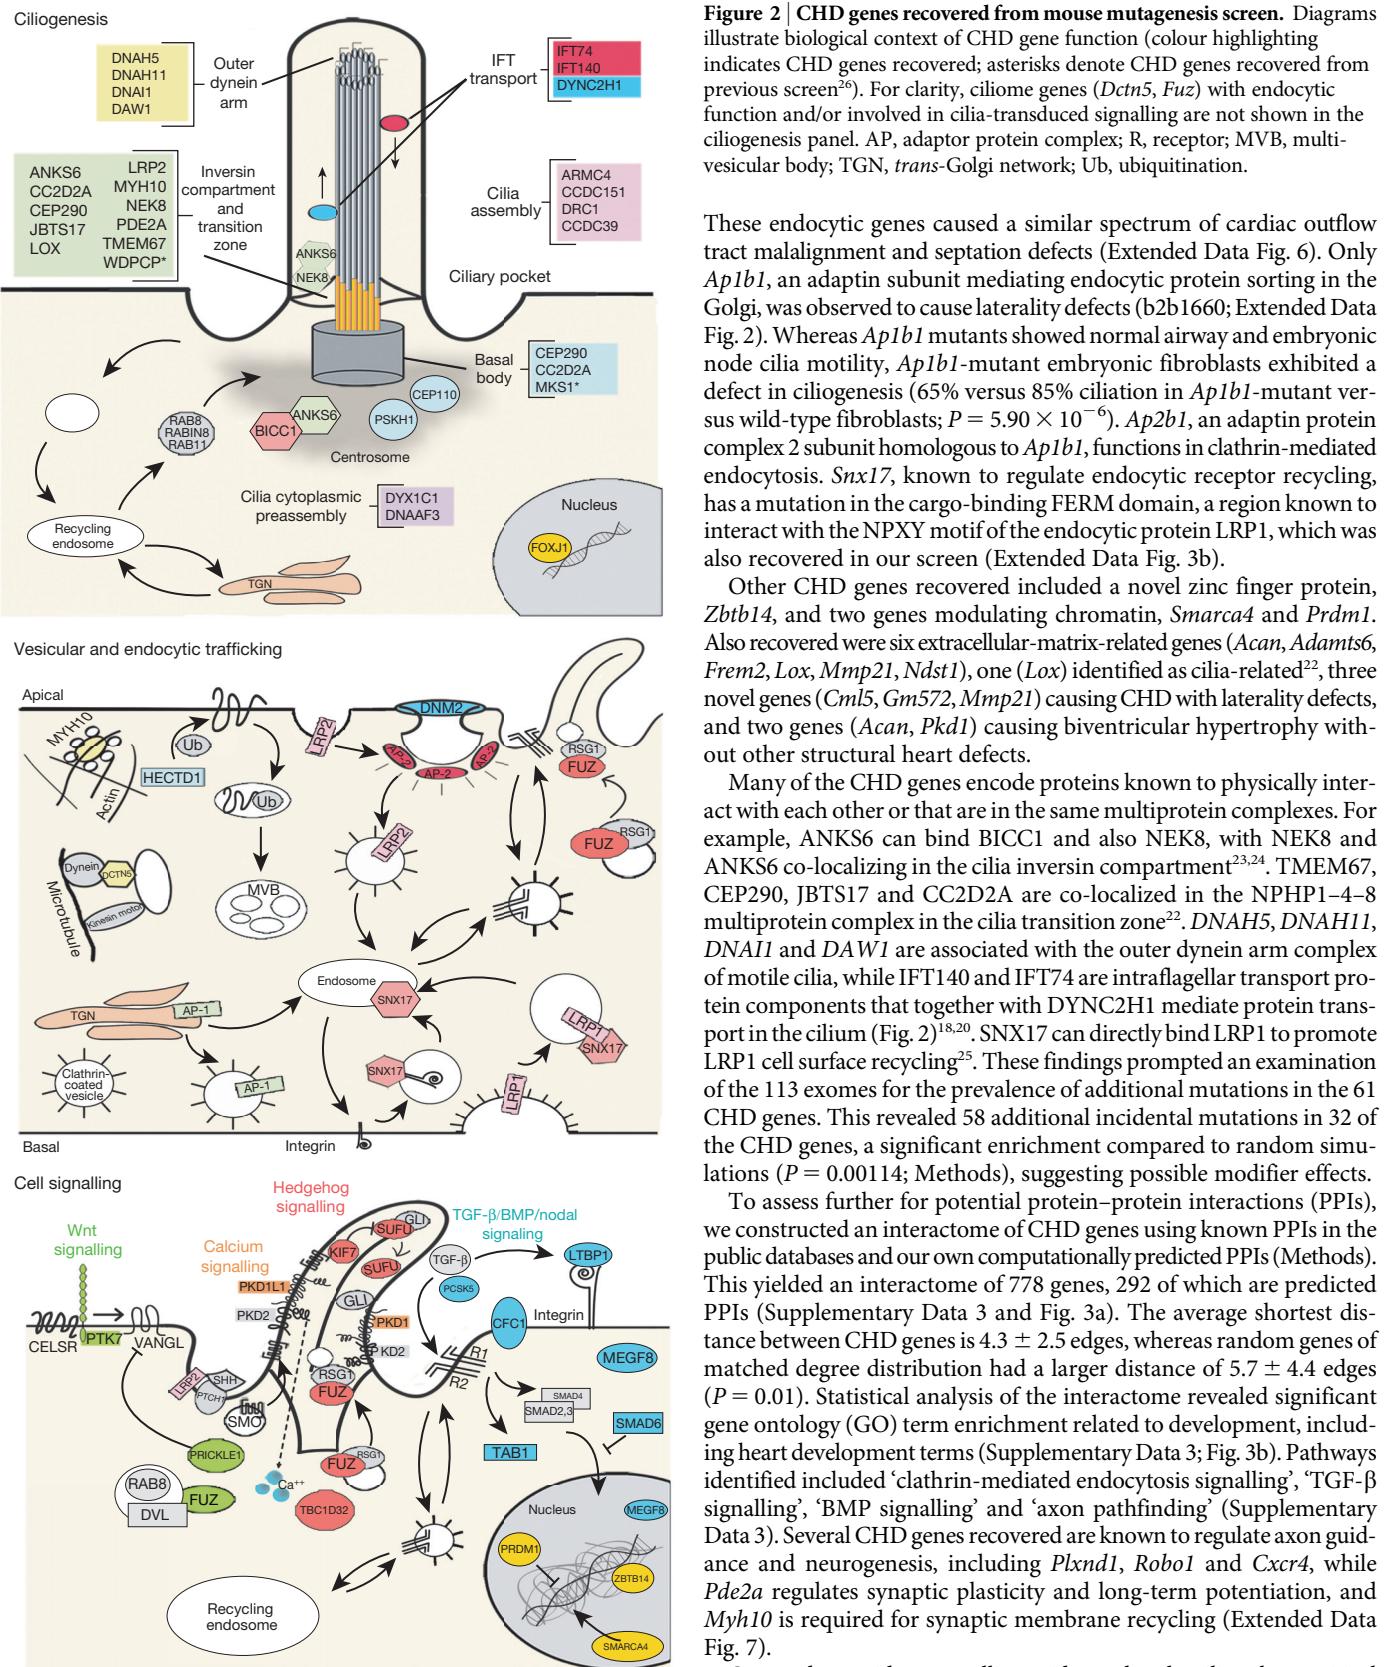


Figure 2 | CHD genes recovered from mouse mutagenesis screen. Diagrams illustrate biological context of CHD gene function (colour highlighting indicates CHD genes recovered; asterisks denote CHD genes recovered from previous screen²⁶). For clarity, ciliome genes (*Dctn5*, *Fuz*) with endocytic function and/or involved in cilia-transduced signalling are not shown in the ciliogenesis panel. AP, adaptor protein complex; R, receptor; MVB, multi-vesicular body; TGN, trans-Golgi network; Ub, ubiquitination.

These endocytic genes caused a similar spectrum of cardiac outflow tract malalignment and septation defects (Extended Data Fig. 6). Only *Ap1b1*, an adaptin subunit mediating endocytic protein sorting in the Golgi, was observed to cause laterality defects (b2b1660; Extended Data Fig. 2). Whereas *Ap1b1* mutants showed normal airway and embryonic node cilia motility, *Ap1b1*-mutant embryonic fibroblasts exhibited a defect in ciliogenesis (65% versus 85% ciliation in *Ap1b1*-mutant versus wild-type fibroblasts; $P = 5.90 \times 10^{-6}$). *Ap2b1*, an adaptin protein complex 2 subunit homologous to *Ap1b1*, functions in clathrin-mediated endocytosis. *Snx17*, known to regulate endocytic receptor recycling, has a mutation in the cargo-binding FERM domain, a region known to interact with the NPXY motif of the endocytic protein LRP1, which was also recovered in our screen (Extended Data Fig. 3b).

Other CHD genes recovered included a novel zinc finger protein, *Zbtb14*, and two genes modulating chromatin, *Smarca4* and *Prdm1*. Also recovered were six extracellular-matrix-related genes (*Acan*, *Adams6*, *Frem2*, *Lox*, *Mmp21*, *Ndst1*), one (*Lox*) identified as cilia-related²², three novel genes (*Cml5*, *Gm572*, *Mmp21*) causing CHD with laterality defects, and two genes (*Acan*, *Pkd1*) causing biventricular hypertrophy without other structural heart defects.

Many of the CHD genes encode proteins known to physically interact with each other or that are in the same multiprotein complexes. For example, *ANKS6* can bind *BICC1* and also *NEK8*, with *NEK8* and *ANKS6* co-localizing in the cilia inversin compartment^{23,24}. *TMEM67*, *CEP290*, *JBTS17* and *CC2D2A* are co-localized in the NPHP1–4–8 multiprotein complex in the cilia transition zone²². *DNAH5*, *DNAH11*, *DNAI1* and *DAW1* are associated with the outer dynein arm complex of motile cilia, while *IFT140* and *IFT74* are intraflagellar transport protein components that together with *DYNC2H1* mediate protein transport in the cilium (Fig. 2)^{18,20}. *SNX17* can directly bind LRP1 to promote LRP1 cell surface recycling²⁵. These findings prompted an examination of the 113 exomes for the prevalence of additional mutations in the 61 CHD genes. This revealed 58 additional incidental mutations in 32 of the CHD genes, a significant enrichment compared to random simulations ($P = 0.00114$; Methods), suggesting possible modifier effects.

To assess further for potential protein–protein interactions (PPIs), we constructed an interactome of CHD genes using known PPIs in the public databases and our own computationally predicted PPIs (Methods). This yielded an interactome of 778 genes, 292 of which are predicted PPIs (Supplementary Data 3 and Fig. 3a). The average shortest distance between CHD genes is 4.3 ± 2.5 edges, whereas random genes of matched degree distribution had a larger distance of 5.7 ± 4.4 edges ($P = 0.01$). Statistical analysis of the interactome revealed significant gene ontology (GO) term enrichment related to development, including heart development terms (Supplementary Data 3; Fig. 3b). Pathways identified included ‘clathrin-mediated endocytosis signalling’, ‘TGF-β signalling’, ‘BMP signalling’ and ‘axon pathfinding’ (Supplementary Data 3). Several CHD genes recovered are known to regulate axon guidance and neurogenesis, including *Plxnd1*, *Robo1* and *Cxcr4*, while *Pde2a* regulates synaptic plasticity and long-term potentiation, and *Myh10* is required for synaptic membrane recycling (Extended Data Fig. 7).

Our study provides compelling evidence that the cilium has a critical role in CHD pathogenesis. The observed enrichment for cilia and cilia-related genes emerged from a phenotype-driven forward genetic screen agnostic to the specific genes that may be recovered. The finding of ciliome mutations in CHD mutant lines with and without laterality defects would suggest a broad role for cilia in CHD pathogenesis. While our study is focused on monogenic recessive coding mutations, we noted

We also recovered ten CHD genes mediating vesicular/endocytic trafficking, including four (*Lrp2*, *Myh10*, *Dctn5*, *Fuz*) that are ciliome-related (Fig. 2). The enrichment for endocytic genes may be mechanistically linked to cilia, as vesicular trafficking at the cilia base has important roles in both ciliogenesis and cilia-transduced cell signalling^{18,21}.

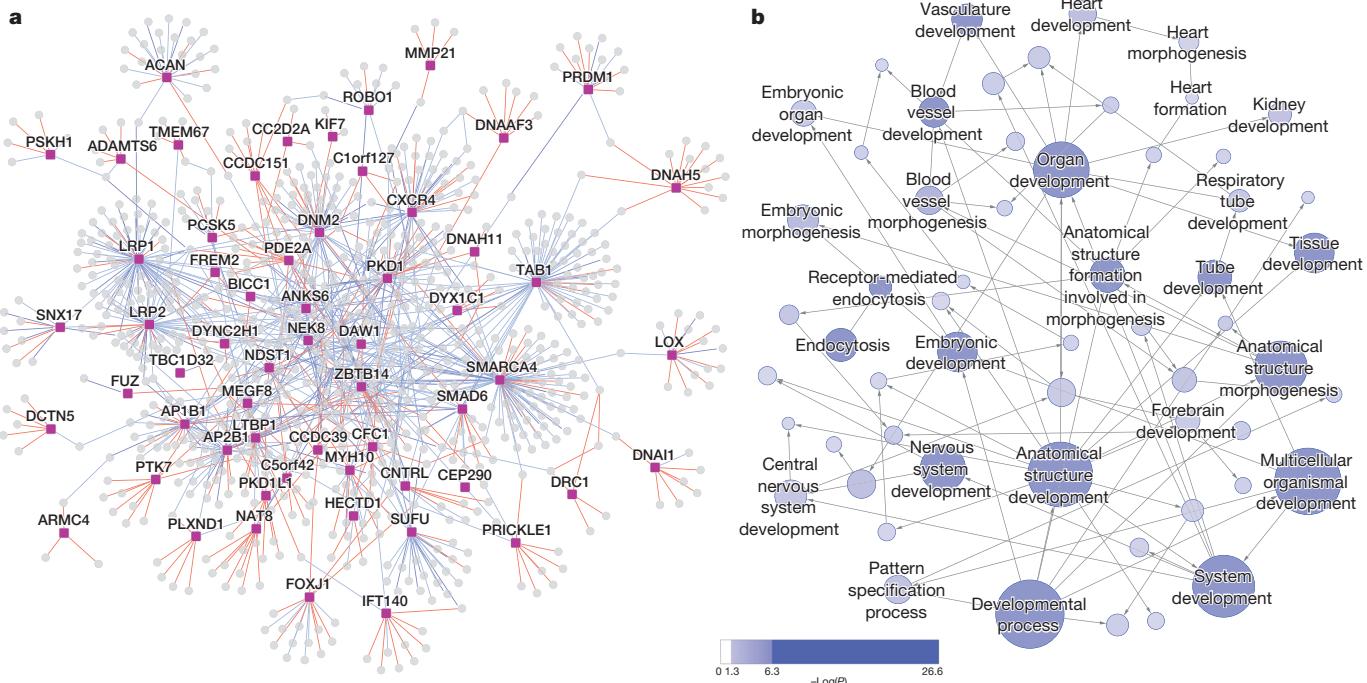


Figure 3 | Interactome network of CHD genes with known and predicted interactors. **a**, An interactome network was constructed comprising the 61 CHD genes (magenta squares) and their known (blue edges) and predicted

a striking overlap of genes recovered in the mouse screen versus heterozygous *de novo* coding mutations in CHD candidate genes identified in CHD patients⁵. Of 28 CHD genes identified in the human study, 11 are cilia-related, 7 are in the TGF- β /SHH/WNT signalling pathway, and 2 are involved in vesicular/endocytic trafficking (Supplementary Data 4). Heterozygous mutation in *PITX2*, a gene essential for left-right patterning¹², was recovered in a non-heterotaxy CHD patient, suggesting a broader role for left-right patterning in CHD pathogenesis. The human study also identified an important role for genes regulating chromatin. Consistent with this, our screen recovered two mouse CHD genes modulating chromatin, *Smarca4* and *Prdm1* (ref. 5). Hence, findings from our fetal mouse screen may provide insights not only into the Mendelian recessive genetic contribution to CHD in the unborn, a segment of the CHD population not well studied, but these findings also may have relevance to the complex genetics of CHD overall. The CHD mouse models generated in this study are invaluable resources for interrogating CHD pathogenesis, with the cryopreserved mutant lines providing access to a large library of mutations, including null mutations for genes without KOMP alleles.

Online Content Methods, along with any additional Extended Data display items and Source Data, are available in the online version of the paper; references unique to these sections appear only in the online paper.

Received 9 September 2014; accepted 26 January 2015.

Published online 25 March 2015.

- Reller, M. D., Strickland, M. J., Riehle-Colarusso, T., Mahle, W. T. & Correa, A. Prevalence of congenital heart defects in metropolitan Atlanta, 1998–2005. *J. Pediatr.* **153**, 807–813 (2008).
- Hoffman, J. I. Incidence of congenital heart disease: II. Prenatal incidence. *Pediatr. Cardiol.* **16**, 103–113 (1995).
- Tennstedt, C., Chaoui, R., Korner, H. & Dietel, M. Spectrum of congenital heart defects and extracardiac malformations associated with chromosomal abnormalities: results of a seven year necropsy study. *Heart* **82**, 34–39 (1999).
- Øyen, N. et al. Recurrence of congenital heart defects in families. *Circulation* **120**, 295–301 (2009).
- Zaidi, S. et al. *De novo* mutations in histone-modifying genes in congenital heart disease. *Nature* **498**, 220–223 (2013).
- Liu, X. et al. Interrogating congenital heart defects with noninvasive fetal echocardiography in a mouse forward genetic screen. *Circ Cardiovasc Imaging* **7**, 31–42 (2014).
- Atkinson, D. E. & Drant, S. Diagnosis of heterotaxy syndrome by fetal echocardiography. *Am. J. Cardiol.* **82**, 1147–1149 (1998).
- Berg, C. et al. Prenatal diagnosis of cardiosplenic syndromes: a 10-year experience. *Ultrasound Obstet. Gynecol.* **22**, 451–459 (2003).
- Taketazu, M., Lougheed, J., Yoo, S. J., Lim, J. S. & Hornberger, L. K. Spectrum of cardiovascular disease, accuracy of diagnosis, and outcome in fetal heterotaxy syndrome. *Am. J. Cardiol.* **97**, 720–724 (2006).
- Tan, S. Y. et al. Heterotaxy and complex structural heart defects in a mutant mouse model of primary ciliary dyskinesia. *J. Clin. Invest.* **117**, 3742–3752 (2007).
- Leigh, M. W. et al. Clinical and genetic aspects of primary ciliary dyskinesia/Kartagener syndrome. *Genet. Med.* **11**, 473–487 (2009).
- Shiratori, H. & Hamada, H. The left-right axis in the mouse: from origin to morphology. *Development* **133**, 2095–2104 (2006).
- Damerla, R. R. et al. Ion torrent sequencing for conducting genome-wide scans for mutation mapping analysis. *Mamm. Genome* **25**, 120–128 (2014).
- Coghill, E. L. et al. A gene-driven approach to the identification of ENU mutants in the mouse. *Nature Genet.* **30**, 255–256 (2002).
- Bultman, S. et al. A *Brg1* null mutation in the mouse reveals functional differences among mammalian SWI/SNF complexes. *Mol. Cell* **6**, 1287–1295 (2000).
- Kasarskis, A., Manova, K. & Anderson, K. V. A phenotype-based screen for embryonic lethal mutations in the mouse. *Proc. Natl. Acad. Sci. USA* **95**, 7485–7490 (1998).
- Chen, J., Knowles, H. J., Hebert, J. L. & Hackett, B. P. Mutation of the mouse hepatocyte nuclear factor/forkhead homologue 4 gene results in an absence of cilia and random left-right asymmetry. *J. Clin. Invest.* **102**, 1077–1082 (1998).
- Sung, C. H. & Leroux, M. R. The roles of evolutionarily conserved functional modules in cilia-related trafficking. *Nature Cell Biol.* **15**, 1387–1397 (2013).
- Ware, S. M., Aygun, M. G. & Hildebrandt, F. Spectrum of clinical diseases caused by disorders of primary cilia. *Proc. Am. Thorac. Soc.* **8**, 444–450 (2011).
- Fliegauf, M., Benzing, T. & Omran, H. When cilia go bad: cilia defects and ciliopathies. *Nature Rev. Mol. Cell Biol.* **8**, 880–893 (2007).
- Clement, C. A. et al. TGF- β signaling is associated with endocytosis at the pocket region of the primary cilium. *Cell Rep.* **3**, 1806–1814 (2013).
- Sang, L. et al. Mapping the NPHP-JBTS-MKS protein network reveals ciliopathy disease genes and pathways. *Cell* **145**, 513–528 (2011).
- Stagner, E. E., Bourette, D. J., Cheng, J. & Bryda, E. C. The polycystic kidney disease-related proteins Bicc1 and SamCystin interact. *Biochem. Biophys. Res. Commun.* **383**, 16–21 (2009).
- Hoff, S. et al. ANKS6 is a central component of a nephronophthisis module linking NEK8 to INVS and NPHP3. *Nature Genet.* **45**, 951–956 (2013).
- Donoso, M. et al. Polarized traffic of LRP1 involves AP1B and SNX17 operating on Y-dependent sorting motifs in different pathways. *Mol. Biol. Cell* **20**, 481–497 (2009).
- Shen, Y. et al. Cardiovascular phenotyping of fetal mice by noninvasive high-frequency ultrasound facilitates recovery of ENU-induced mutations causing congenital cardiac and extracardiac defects. *Physiol. Genomics* **24**, 23–36 (2005).

Supplementary Information is available in the online version of the paper.

Acknowledgements We thank R. Ramirez for early assistance with the mutagenesis breeding pipeline, R. Subramanian and D. Farkas for early assistance with necropsy and pathology examination of mutants, A. Srinivasan for early assistance with exome sequencing, S. Fataki for assistance with sequencing data maintenance, M. Wong and C. Krise for assistance with mouse curation, B. Beutler for advice on mapping mutations using intercrosses with the C57BL/10J strain and whole-mouse exome sequencing analysis, D. Weeks and Y. Shan for assistance in statistical modelling of target gene size estimates, E. Goldmuntz for helpful discussions and critical review of the manuscript, and the New England Research Institutes (NERI) for constructing the CHD Mouse Mutation Database. The project was supported by award numbers U01HL098180 (to C.W.L.) and U01HL098188 (to NERI) from the National Heart, Lung, and Blood Institute, R01MH094564 (to M.K.G.) from the National Institute of Mental Health, and HG000330 (to J.E.) from the National Human Genome Research Institute. Funding was also provided by the University of Pittsburgh School of Medicine. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Heart, Lung, and Blood Institute, the National Human Genome Research Institute or the National Institutes of Health.

Author Contributions Study design: C.W.L. ENU mutagenesis, line cryopreservation and JAX strain datasheet construction: H.P., L.R., J.M., L.H. Mouse breeding, sample collection, sample tracking: S.A., C.L., K.L., G.C.G., A.V., C.W.L. Electronic database construction and maintenance: C.C. MGI curation: K.T., G.C.G., L.L., C.W.L., C.L.S., J.T.E CHD phenotyping: X.L., K.L., Y.C., G.C.G., A.J.K., S.A., W.D., C.W.L., L.L., K.T., R.F. Cilia

immunostain and histology: J.T.S.A., G.J.P., R.F. Analysis of airway and node cilia motility: R.F., K.L., G.C.G., A.J.K. Exome sequencing analysis: Y.L. Mutation validation: N.T.K., B.C., R.R.D., H.Y., Y.L. Mutation mapping: R.R.D., N.T.K., B.C., Y.L. Interactome analysis: M.K.G., M.T. Ciliome and pathway annotation: C.W.L., G.J.P., G.C.G., N.T.K., Y.L. Manuscript preparation: C.W.L., Y.L., N.T.K., G.C.G.

Author Information All mutant mouse lines recovered in this mouse mutagenesis screen and their phenotype description and causative mutations are curated in the MGI database (<http://www.informatics.jax.org>) and can be retrieved by entering “b2b” in the search box. All mutant mouse lines curated in MGI can be reanimated from sperm cryopreserved in the Jackson Laboratory (JAXMice) Repository. All mutations recovered by mouse exome sequencing analysis are searchable together with phenotype information via the public Bench to Bassinet Congenital Heart Disease Mouse Mutation Database (<http://benchtobassinet.com/ForResearchers/BasicScienceDataResourceSharing/GeneDiscoveryinMouseModels.aspx>). The mouse exome datasets are available from the GNOMEx Cardiovascular Development Consortium Datahub (<https://b2b.hci.utah.edu/gnomex/gnomexGuestFlex.jsp?topicNumber=67>). Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to C.W.L. (cel36@pitt.edu).

METHODS

CHD phenotyping. C57BL/6J mice were chemically mutagenized using ethylnitrosourea and subject to a two generation backcross breeding scheme²⁶. To identify CHD mutants, G3 fetuses were ultrasound scanned with a two-tier ultrasound phenotyping strategy using the Acuson Sequoia and Vevo2100 ultrasound systems as previously described⁶. The fetal ultrasound imaging was conducted at E13.5–15.5, after completion of major cardiovascular developmental processes, including outflow tract septation and ventricular chamber septation. All of the CHD diagnoses were confirmed or refined by additional analysis using micro-CT/micro-MRI analyses and three-dimensional histopathology using episcopic confocal microscopy^{6,27} (Extended Data Fig. 1). The summative CHD diagnosis and phenotype description are curated in the public Mouse Genome Informatics (MGI) database (<http://www.informatics.jax.org>; enter “b2b” into the search box to pull up all the lines or type in specific line IDs or MGI IDs to go directly to a specific mutant line). Included in each mutant line MGI record are sample images and videos collected from various different imaging modalities, including ultrasound images and videos, necropsy images, micro-CT/micro-MRI scans, episcopic confocal microscopy image stacks and three-dimensional reconstructions, histology, and high-speed videos for analysis of tracheal airway cilia motility. All animal experiments described in this study were conducted with approved Institutional Animal Care and Use Committee (IACUC) protocols from the University of Pittsburgh.

Cilia videomicroscopy. Trachea of newborn mice were harvested and mounted as previously described. For nodal cilia imaging, the node-containing region of E7.5–E7.75 embryos (3–6 somite stages) was dissected and mounted on a glass cover-slip¹⁰. Videomicroscopy was carried out using a Leica DMIRE2 inverted microscope at 37 °C with a $\times 100$ oil objective, and DIC optics. To examine flow, 0.35 μm microspheres were added to the medium bathing the ciliated tissue samples ($\sim 1 \times 10^{11}$ particles per ml; Polysciences). High-speed videos (200 fps) were collected using a Phantom v.4.2 camera (Vision Research).

Cilia immunostaining analysis. E12.5 mutant and control littermate embryos from b2b1035 were immunostained with ciliary protein antibodies including acetylated α -tubulin antibody (Sigma T7451, 1:1,000) and IFT88 (1:500), as previously described²⁸.

Mouse embryonic fibroblasts derived from E11.5–E12.5 b2b1660 mutant and control littermate embryos were serum starved for 24 to 48 h to grow out the cilia, with staining of ciliary proteins carried out using primary antibodies antibodies against acetylated α -tubulin (Sigma T7451, 1:1,000) and γ -tubulin antibody (Sigma T6557, 1:1,000). All immunostaining samples were imaged by confocal microscopy.

Mutation recovery with whole-exome sequencing analysis. Genomic DNA from CHD mutant mice were sequence captured using Agilent SureSelect Mouse All Exon kit V1, and sequenced using Illumina HiSeq 2000 with minimum average 50 \times target sequence coverage (BGI Americas). Sequence reads were aligned to the C57BL/6 mouse reference genome (mm9) and analysed using CLC Bio Genomic Workbench and GATK software. To minimize false negatives, variant calls were set at 5 \times minimum coverage and $\geq 20\%$ alternate reads. Sequence variants were annotated with annoVar (<http://www.openbioinformatics.org/annovar/>) and filtered against dbSNP128 and our in-house mouse exome databases with custom scripts. Intronic variants within 10 bp from the splice junction were annotated as potential splicing mutation. All of the mutations recovered are searchable together with phenotype information via the public Bench to Bassinet Congenital Heart Disease Mouse Mutation Database (<http://benctobassinet.com/ForResearchers/BasicScienceDataResourceSharing/GeneDiscoveryinMouseModels.aspx>). The exome data sets are available from the GNOMEx CvDC Datahub (<https://b2b.hci.utah.edu/gnomex/gnomexGuestFlex.jsp?topicNumber=67>).

Ciliome gene annotation. Ciliome genes were derived from a compilation of multiple bioinformatics, genomic and proteomic studies in different species²⁹. This included a bioinformatics study of X-box-containing genes in *Caenorhabditis elegans* and several proteomic and comparative genomic studies comprising ciliated organisms in several species (*C. elegans*, *Drosophila*, human, *Chlamydomonas*, *Plasmodium*,

falciparum and *Trypanosoma*). In addition, we also included genes recently identified in the cilia transition zone complex²² and genes identified to cause primary ciliary dyskinesia ([http://ghr.nlm.nih.gov/condition/primary-ciliary-dyskinesia/show/Related+Gene\(s\)](http://ghr.nlm.nih.gov/condition/primary-ciliary-dyskinesia/show/Related+Gene(s))).

Analysis of splicing defect mutations. For analysis of splicing defect mutations, RNA was isolated from skin, trachea or heart tissue and reverse transcribed to generate cDNA using a High Capacity RNA-to-cDNA Kit (Life Technologies). PCR amplification of the cDNA was carried out with primers spanning the affected splice junction using AmpliTaq Gold DNA polymerase system (Life Technologies) (95 °C for 5 min, 40 cycles at 95 °C for 30 s, 58 °C for 30 s at 72 °C for 1 min, 72 °C for 5 min). The amplified products were separated on a 3% agarose gel and the bands were excised using the QIAquick Gel Extraction Kit (Qiagen) and analysed by Sanger capillary sequencing.

Simulation-based mutation enrichment analysis. A simulation approach was used to test whether mutations in the 61 CHD genes are enriched in the mouse exome data sets generated from the 113 CHD mutant mouse lines. From the 113 exomes, we recovered 149 mutations in the 61 CHD genes, 91 of which were identified as pathogenic, yielding an excess of 58 incidental mutations. For this simulation, all the genes from the mouse genome were grouped in bins of 500 genes based on coding DNA sequence (CDS) length. Then 61 genes with the same size distribution as the observed 61 CHD genes were randomly sampled. For each bin, we randomly sampled the same number of genes as those observed for the CHD genes, and counted the number of mutations in these genes in the 113 mouse exomes. On the basis of 50,000 sets of 61 randomly sampled genes, we calculated a *P*-value based on the count of random gene sets that harboured 149 or more mutations. To take potential gene coverage bias into account, we limited our random gene selection to genes with at least one coding mutation in our CHD exomes, which yielded a *P* value of 0.00114. When we choose random genes from all genes, including genes not found in our exome data set, the *P* value is <0.00002.

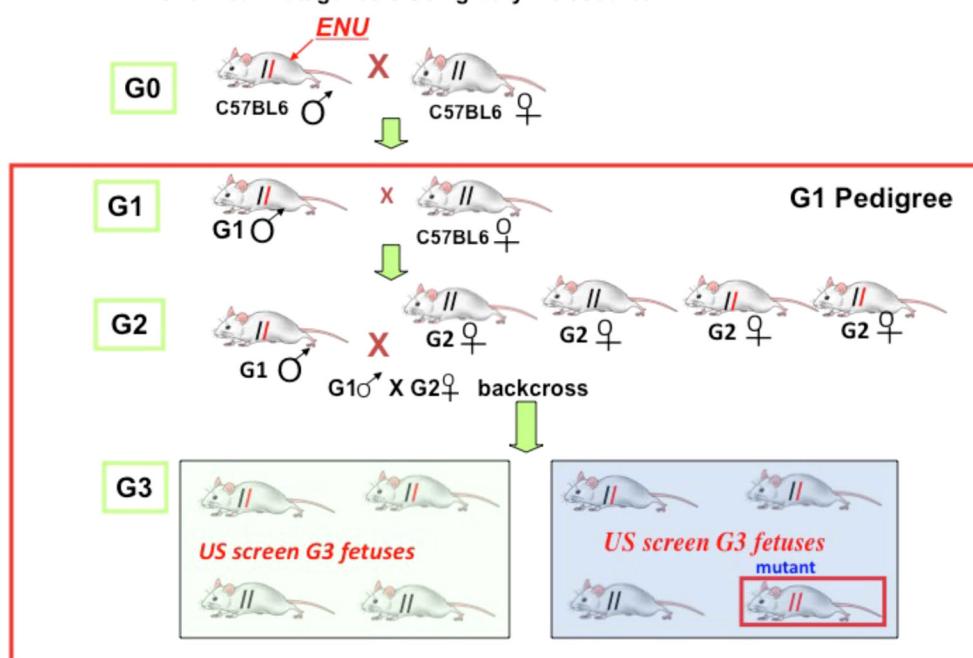
Interactome analysis. To construct the interactome, each of the mouse CHD genes was mapped to its human orthologue, and PPIs of these genes were extracted from the Human Protein Reference Database (HPRD; <http://www.hprd.org>) and BioGRID (<http://thebiogrid.org>). The interactome is augmented with PPIs discovered computationally using the recently developed HiPIP model (unpublished results). Pathway and GO term enrichment analyses of the interactome were carried out using Ingenuity Pathway Analysis suite and BiNGO³⁰ plugin of Cytoscape (<http://www.cytoscape.org>). *P* value for each pathway and GO term was calculated using Fisher's exact test with Benjamini–Hochberg correction (Supplementary Data 3).

Shortest paths were computed using Dijkstra's algorithm implemented in the networkx package of python. Average length of the shortest paths between pairs of CHD genes was compared to that of 61 random genes, where random genes were sampled to have similar degree distribution as the 61 CHD genes. For this, all human genes were categorized into 7 bins based on their degree in the interactome (0, 1–5, 6–25, 26–50, 51–75, 75–100, >100). As many random genes were selected from any bin as the number of CHD genes in that bin. This process was repeated 1,500 times, and average length of the shortest paths was computed each time. *P* value is computed empirically to determine statistical significance.

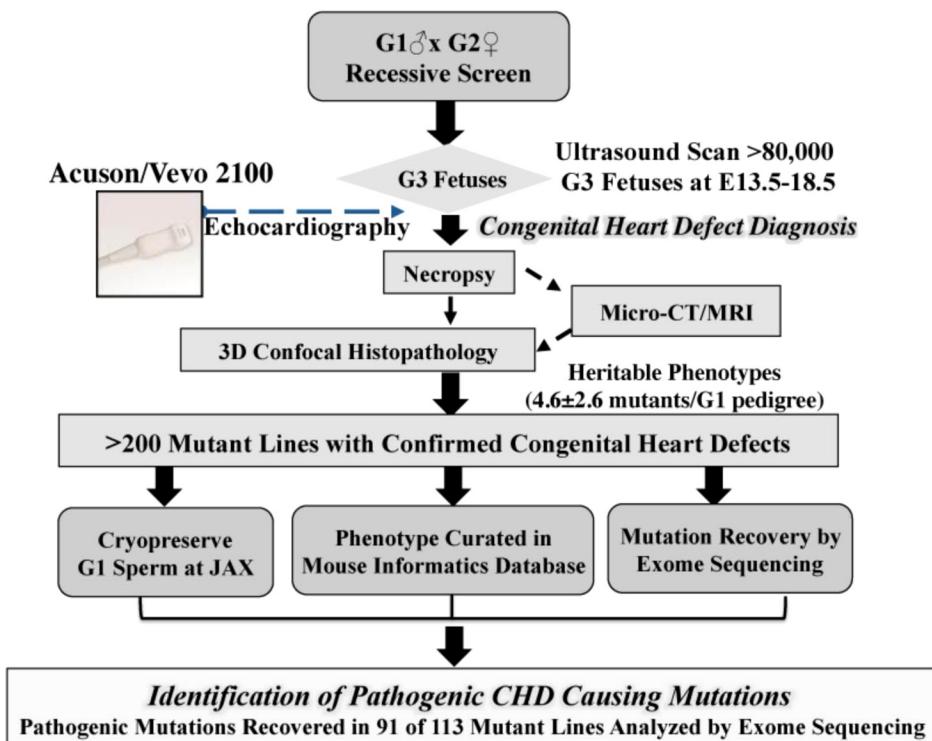
27. Kim, A. J. *et al.* Microcomputed tomography provides high accuracy congenital heart disease diagnosis in neonatal and fetal mice. *Circ Cardiovasc Imaging* **6**, 551–559 (2013).
28. Pazour, G. J. *et al.* The intraflagellar transport protein, IFT88, is essential for vertebrate photoreceptor assembly and maintenance. *J. Cell Biol.* **157**, 103–114 (2002).
29. Inglis, P. N., Boroevich, K. A. & Leroux, M. R. Piecing together a ciliome. *Trends Genet.* **22**, 491–500 (2006).
30. Maere, S., Heymans, K. & Kuiper, M. BiNGO: a Cytoscape plugin to assess overrepresentation of gene ontology categories in biological networks. *Bioinformatics* **21**, 3448–3449 (2005).

a Two Generation Backcross Breeding Scheme

Chemical Mutagenesis Using Ethylnitrosourea

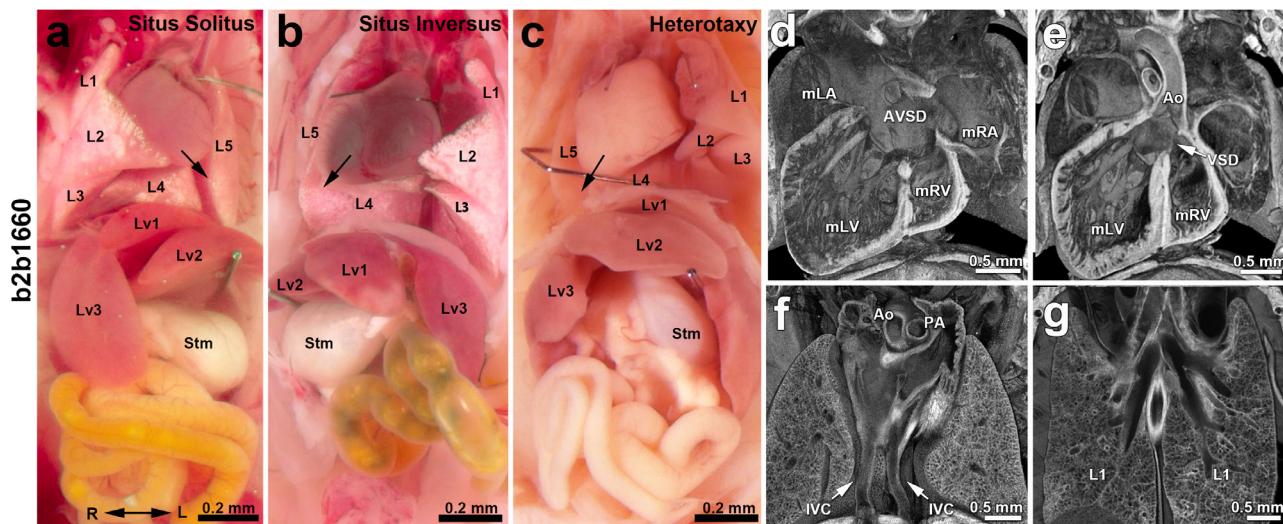


b Mouse Forward Genetic Screen Workflow



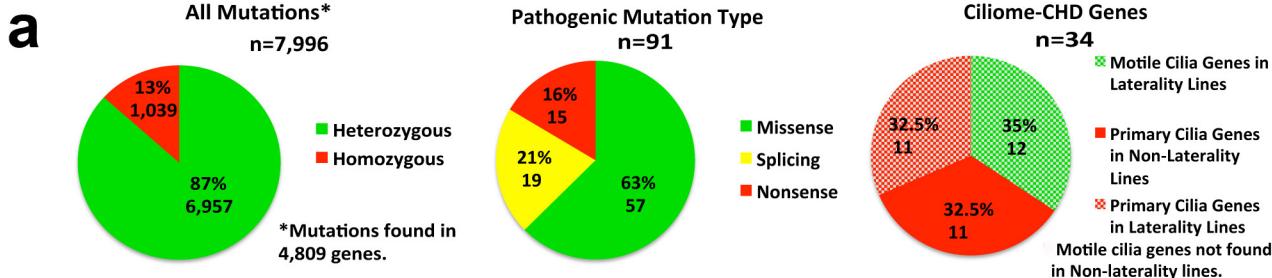
Extended Data Figure 1 | Breeding, phenotyping and mutation recovery pipeline for mouse forward genetic screen. **a**, Two generation backcross breeding scheme used to generate G3 mutants with recessive mutations causing congenital heart defects, with all offspring from a single G1 male defined

as a distinct pedigree or mutant line. **b**, Pipeline for recovery and curation and cryopreservation of CHD mutant mouse models and the recovery of pathogenic CHD causing mutations.



Extended Data Figure 2 | Situs anomalies and congenital heart defects in *Ap1b1^{b2b1660}* mutants. **a–c**, Mutants from line 1660, identified with an *Ap1b1* mutation, exhibit situs solitus (**a**), situs inversus (**b**) or heterotaxy (**c**). Situs solitus, characterized by normal left–right visceral organ positioning, the heart apex (arrow) points to the left (levocardia), four lung lobes are on the right and one on the left, stomach is to the left, and the dominant liver lobe is on the right. With situs inversus, there is complete mirror reversal of organ situs, while with heterotaxy, visceral organ situs is randomized, such as dextrocardia

with levogastria shown in **c**. **d–g**, The heterotaxy mutant in **c** exhibits complex CHD with atrioventricular septal defect (AVSD) (**d**), ventricular septal defect (VSD) (**e**), duplicated inferior vena cava (IVC) (**f**) and left pulmonary isomerism with bilateral single lung lobes (**g**). Ao, aorta; L1–5, lung lobes 1–5; Lv1–3, liver lobes 1–3; mLA, morphologic left atrium; mLV, morphologic left ventricle; mRA, morphologic right atrium; mRV, morphologic right ventricle; PA, pulmonary artery; Stm, stomach.



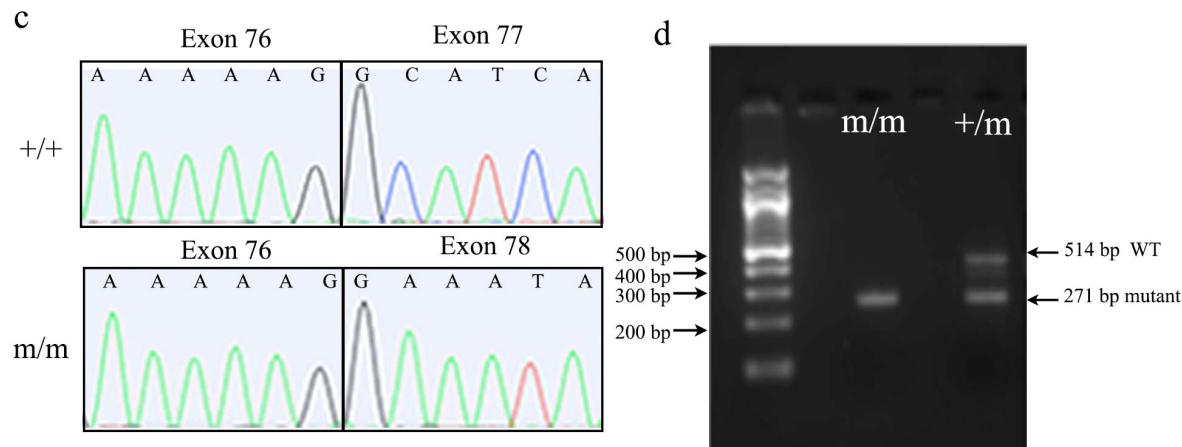
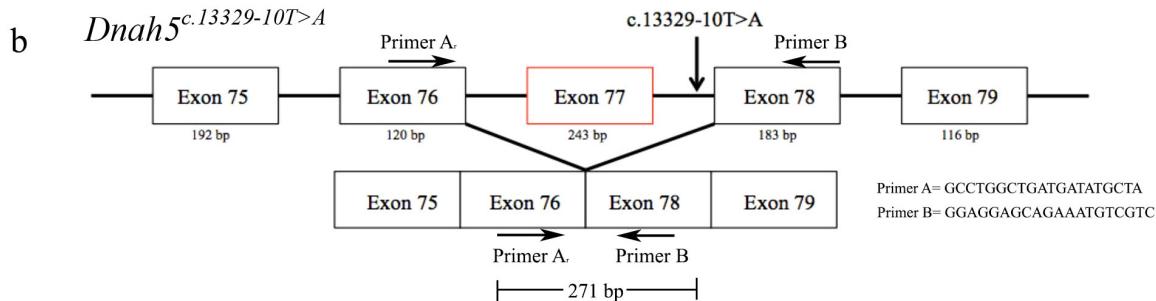
b Recovery of Pathogenic Mutations

Gene	Cardiac and Laterality Phenotype	Gene	Cardiac Phenotype
<i>Anks6</i>	Complex CHD with heterotaxy	<i>Acan</i>	Biventricular hypertrophy
<i>Ap1b1</i>	Complex CHD with heterotaxy	<i>Adamts6</i>	DORV, overriding aorta, AVSD, ventricular noncompaction
<i>Armc4</i>	Complex CHD with heterotaxy	<i>Ap2b1</i>	TGA,AVSD,VSD, aortic arch defects, biventricular hypertrophy
<i>Bicc1</i>	Complex CHD with heterotaxy	<i>Cep110</i>	DORV, AVSD, VSD, RAA, hypoplastic PA
<i>Cc2d2a</i>	Complex CHD with heterotaxy	<i>Cxcr4</i>	Ao atresia,hypoplastic transverse arch,double Ao arch, VSD
<i>Ccdc151</i>	Complex CHD with heterotaxy	<i>Dctn5</i>	DORV, overriding aorta, VSD
<i>Ccdc39</i>	Complex CHD with heterotaxy	<i>Dnm2</i>	DORV, overriding Ao, hypoplastic Ao, aortic arch defects,VSD
<i>Cep290</i>	Complex CHD with heterotaxy	<i>Frem2</i>	Biventricular hypertrophy
<i>Cfc1</i>	TGA with heterotaxy	<i>Fuz</i>	PTA, AVSD, RAA, MAPCA
<i>Cml5</i>	Complex CHD with heterotaxy	<i>Hectd1</i>	Aortic atresia,hypoplastic Ao,dysplastic semilunar valves,VSD
<i>Daw1</i>	Complex CHD with heterotaxy	<i>Jbis17</i>	DORV, pulmonary stenosis/atresia, AVSD
<i>Dnaaf3</i>	Complex CHD with heterotaxy	<i>Kif7</i>	DORV,overriding aorta, hypoplastic PA, AVSD, RAA/IAA
<i>Dnah11</i>	Complex CHD with heterotaxy	<i>Lox</i>	Aortic stenosis/thickened aorta, biventricular hypertrophy
<i>Dnah5</i>	Complex CHD with heterotaxy	<i>Lrp1</i>	DORV, pulmonary stenosis, PTA, AVSD
<i>Dnai1</i>	Complex CHD with heterotaxy	<i>Lrp2</i>	PTA, VSD, aortic arch defects, ventricular noncompaction
<i>Drc1</i>	Complex CHD with heterotaxy	<i>Ltbp1</i>	DORV, PTA, VSD, aortic arch defects
<i>Dync2h1</i>	Laterality defects with CHD, MAPCA	<i>Myh10</i>	DORV,overriding Ao,AVSD, IAA,ventricular noncompaction
<i>Dyx1c1</i>	Complex CHD with heterotaxy	<i>Ndst1</i>	PTA, Ao arch defects, AVSD, biventricular hyperrophy
<i>Foxj1</i>	Complex CHD with heterotaxy	<i>Pde2a</i>	DORV,overriding aorta,VSD,AVSD, dys SL valves, ventricular NC
<i>Gm572</i>	Complex CHD with heterotaxy	<i>Pkd1</i>	Biventricular hypertrophy
<i>Ift140</i>	Laterality defects with CHD	<i>Plxnd1</i>	PTA,DORV,overriding Ao, hypoplastic PA,AVSD, RAA
<i>Ift74</i>	Laterality defects with CHD	<i>Prdm1</i>	DORV, AVSD, aortic valve atresia, VSD, RAA/vascular ring
<i>Megf8</i>	TGA with heterotaxy	<i>Pickle1</i>	DORV, overriding aorta, hypoplastic PA
<i>Mmp21</i>	TGA with heterotaxy	<i>Ptk7</i>	DORV, VSD, biventricular hypertrophy
<i>Nek8</i>	Complex CHD with heterotaxy	<i>Robo1</i>	DORV, VSD, biventricular hypertrophy
<i>Pcsk5</i>	Dextro/mesocardia,PTA/DORV/AVSD MAPCA	<i>Smad6</i>	PTA,DORV,TGA,Ao valve dysplasia,hypoplastic PA,RAA/IAA,VSD
<i>Pkd1l1</i>	Complex CHD with heterotaxy	<i>Smarca4</i>	DORV,overriding aorta VSD
<i>Pskh1</i>	Complex CHD with heterotaxy	<i>Snx17</i>	DORV/overriding aorta, AVSD
<i>Tbc1d32</i>	Complex CHD with heterotaxy	<i>Sufu</i>	Overriding aorta, VSD, coronary artery defects
<i>Tmem67</i>	Complex CHD with heterotaxy	<i>Tab1</i>	PTA,DORV,overriding Ao,AVSD,hypoplastic transv arch/IAA/VS
		<i>Zbtb14</i>	DORV, AVSD, semilunar valve defects

Extended Data Figure 3 | Distribution of pathogenic mutations recovered from the forward genetic screen. a, Distribution of all incidental coding mutations (left), pathogenic mutations (middle) and ciliome CHD genes (right) recovered from 113 CHD mouse mutant lines. **b,** Recovery of pathogenic mutations and associated CHD phenotypes. Grey-filled boxes indicate CHD mutations in genes not previously identified to cause CHD. Ao, aorta; AVSD,

atrioventricular septal defect; BVH, biventricular hypertrophy; DORV, double outlet right ventricle; IAA, interrupted aortic arch; MAPCA, major aortopulmonary collateral artery; PA, pulmonary artery; PTA, persistent truncus arteriosus; RAA, right aortic arch; TGA, transposition of the great arteries; VSD, ventricular septal defect; VS, vascular sling.

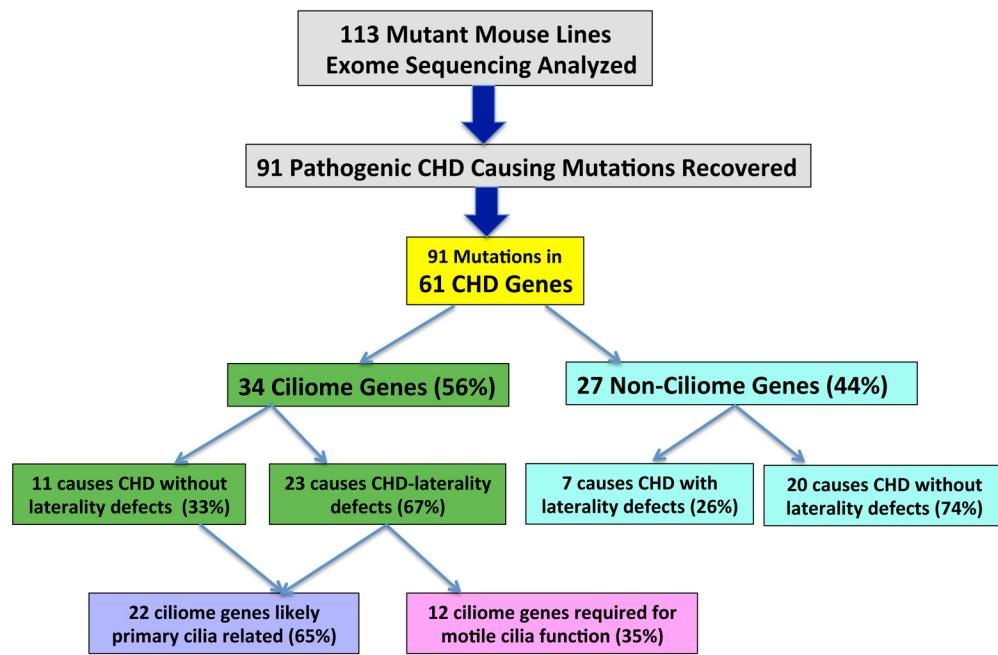
a	Gene	Chr	Position	Nucleotide Change	Bases from exon (bp)	Splicing Site Affected
	<i>Dnai1</i>	4	41544973	c.240+1G>A	1	Donor
	<i>Gm572</i>	4	148017493	c.26+1G>A	1	Donor
	<i>Plxnd1</i>	6	115919325	c.2691+1G>A	1	Donor
	<i>Ccdc151</i>	9	21799255	c.828+2T>C	2	Donor
	<i>Ccdc39</i>	3	33737998	c.357+2T>A	2	Donor
	<i>Cep290</i>	10	99962906	c.1189+2T>C	2	Donor
	<i>Dnah11</i>	12	119425356	c.1137+2T>C	2	Donor
	<i>Dnah11</i>	12	119282224	c.6489+2T>C	2	Donor
	<i>Dnah5</i>	15	28275670	c.7398+2T>A	2	Donor
	<i>Dnah5</i>	15	28159416	c.438+2T>A	2	Donor
	<i>Fuz</i>	7	52152701	c.387+2T>A	2	Donor
	<i>Tbc1d32</i>	10	55915160	c.1071+2T>A	2	Donor
	<i>Bicc1</i>	10	70420407	c.606+2T>C	2	Donor
	<i>Ndst1</i>	18	60850830	c.2426+2T>A	2	Donor
	<i>Lrp2</i>	2	69316757	c.8456-3A>G	3	Acceptor
	<i>Tmem67</i>	4	11974951	c.2322+6T>C	6	Donor
	<i>Dnm2</i>	9	21282798	c.1196+6T>A	6	Donor
	<i>Dctn5</i>	7	129281879	c.348+6T>C	6	Donor
	<i>Dnah5</i>	15	28388293	c.13329-10T>A	10	Acceptor



Extended Data Figure 4 | Pathogenic splicing mutations causing CHD.

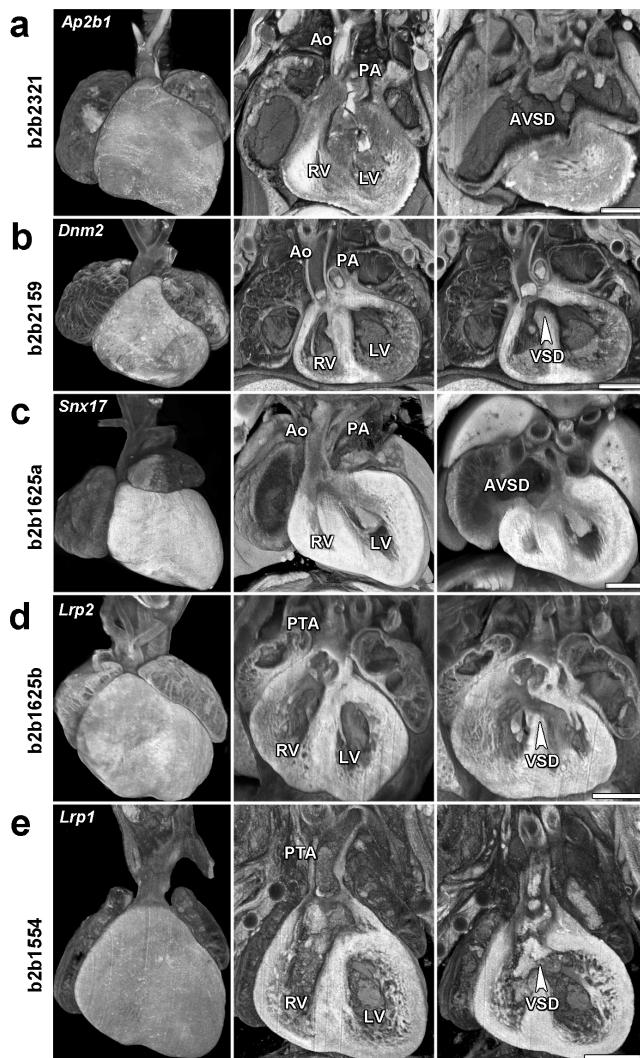
a, The 19 pathogenic splicing mutations recovered are shown, with mutations located beyond the 2-base canonical splice junction highlighted in grey.
b, Schematic diagram showing the anomalous *Dnah5*^{c.13329-10T>A} mutant transcript observed, with the polymerase chain reaction (PCR) primer location and anomalous PCR product size indicated. **c**, Sanger sequencing profile

showing point mutation in *Dnah5*^{c.13329-10T>A} mutant transcript versus that of wild type. **d**, PCR amplification of *Dnah5*^{c.13329-10T>A} heterozygous mutant (m/+) showed the expected 514 bp wild-type and 271 bp mutant PCR product, while only the 271 bp mutant product was observed in the homozygous mutant (m/m) sample.



Extended Data Figure 5 | Ciliome mutations causing CHD with and without laterality defects. A flow chart showing the distribution of ciliome versus non-ciliome CHD genes among laterality versus non-laterality CHD

lines, and further stratification of ciliome CHD genes affecting primary versus motile cilia function.



Extended Data Figure 6 | CHD phenotypes associated with mutations affecting endocytic trafficking. **a–c,** Outflow tract malalignment defects with double outlet right ventricle (DORV) and overriding aorta (Ao) were observed in *Ap2b1* (**a**), *Dnm2* (**b**) and *Snx17* (**c**) mutants, with *Ap2b1* (**a**) mutant also showing anterior positioning of the aorta (Taussig–Bing type DORV). *Snx17* mutant also has AVSD. **d, e,** *Lrp2* (**d**) and *Lrp1* (**e**) mutants both exhibited outflow tract septation defect with persistent truncus arteriosus (PTA). Scale bar, 0.5 mm.

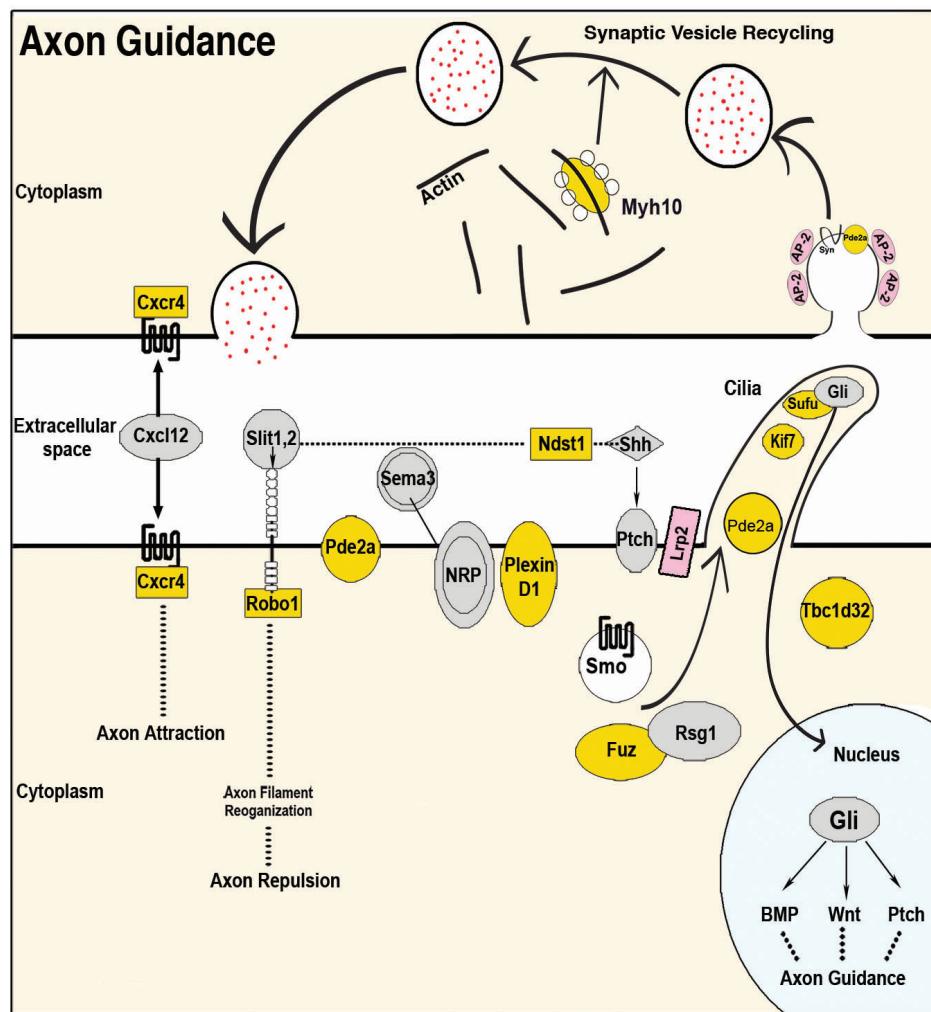

Extended Data Figure 7 | CHD genes associated with axon guidance.

Diagram illustrating the biological context of several CHD genes known to be involved in axonal guidance (colour highlighting indicates CHD genes

recovered from the present screen). Adapted from QIAGEN's Ingenuity Pathway Analysis (<http://www.qiagen.com/ingenuity>).

Extended Data Table 1 | CHD genes with multiple alleles

Gene Symbol (mouse)	Gene Symbol (human)	Mouse CDS (bp)	# Alleles Recovered
<i>Daw1</i>	<i>DAW1</i>	933	2
<i>Dnai1</i>	<i>DNAI1</i>	2,106	2
<i>Drc1</i>	<i>DRC1</i>	2,262	2
<i>Pde2a</i>	<i>PDE2A</i>	2,808	2
<i>Ccdc39</i>	<i>CCDC39</i>	2,814	3
<i>Tmem67</i>	<i>TMEM67</i>	2,988	2
<i>Armc4</i>	<i>ARMC4</i>	3,114	1 [*]
<i>Adamts6</i>	<i>ADAMTS6</i>	3,351	1 [†]
<i>Pcsk5</i>	<i>PCSK5</i>	5,634	2
<i>Plxnd1</i>	<i>PLXND1</i>	5,778	2
<i>Cep290</i>	<i>CEP290</i>	7,440	2
<i>Pkd1l1</i>	<i>PKD1L1</i>	7,824	2
<i>Megf8</i>	<i>MEGF8</i>	8,370	2
<i>Dnah11</i>	<i>DNAH11</i>	13,467	8
<i>Dnah5</i>	<i>DNAH5</i>	13,866	9

* Identical mutation was recovered from two different mutant lines.

† Identical mutation was recovered from three different mutant lines, later confirmed to be a spontaneous mutation in the support colony.

Extended Data Table 2 | Ciliome mutations in laterality and non-laterality lines

Laterality Mutant Lines (30)		Non-Laterality Mutant Lines (31)		
Ciliome Genes	Non-Ciliome Genes	Ciliome Genes	Non-Ciliome Genes	
<i>Anks6</i>	<i>Ap1b1</i>	<i>Cep110</i>	<i>Acan</i>	
<i>Armc4 (m)</i>	<i>Cfc1</i>	<i>Dctn5</i>	<i>Adamts6</i>	
<i>Bicc1</i>	<i>Cml5</i>	<i>Fuz</i>	<i>Ap2b1</i>	
<i>Cc2d2a</i>	<i>Gm572</i>	<i>Jbts17</i>	<i>Cxcr4</i>	
<i>Ccdc151 (m)</i>	<i>Megf8</i>	<i>Kif7</i>	<i>Dnm2</i>	
<i>Ccdc39 (m)</i>	<i>Mmp21</i>	<i>Lrp2</i>	<i>Frem2</i>	
<i>Cep290</i>	<i>Pcsk5</i>	<i>Lox</i>	<i>Hectd1</i>	
<i>Daw1 (m)</i>	Total: 7		<i>Lrp1</i>	
<i>Dnaaf3 (m)</i>		<i>Pde2a</i>	<i>Ltbp1</i>	
<i>Dnah5 (m)</i>		<i>Pkd1</i>	<i>Ndst1</i>	
<i>Dnah11(m)</i>		<i>Sufu</i>	<i>Plxnd1</i>	
<i>Dnai1 (m)</i>		Total: 11		
<i>Drc1 (m)</i>			<i>Prdm1</i>	
<i>Dync2h1</i>			<i>Prickle1</i>	
<i>Dyx1c1 (m)</i>			<i>Ptk7</i>	
<i>Foxj1 (m)</i>			<i>Robo1</i>	
<i>Ifi140</i>			<i>Smad6</i>	
<i>Ifi74</i>			<i>Smarca4</i>	
<i>Nek8</i>			<i>Snx17</i>	
<i>Pkd1l1 (m)</i>			<i>Tab1</i>	
<i>Pskh1</i>			<i>Zbtb14</i>	
<i>Tbc1d32</i>			Total: 20	
<i>Tmem67</i>				
Total: 23 (12)				

m, mutation affects motile cilia function.

Extended Data Table 3 | Mouse CHD genes and associated human diseases

Ciliopathy Related	
<i>Armc4</i>	Primary ciliary dyskinesia, Kartagener Syndrome
<i>Ccdc39</i>	Primary ciliary dyskinesia, Kartagener Syndrome
<i>Dnaaf3</i>	Primary ciliary dyskinesia, Kartagener Syndrome
<i>Dnah11</i>	Primary ciliary dyskinesia, Kartagener Syndrome
<i>Dnah5</i>	Primary ciliary dyskinesia, Kartagener syndrome
<i>Dnai1</i>	Primary ciliary dyskinesia, Kartagener syndrome
<i>Drc1</i>	Primary ciliary dyskinesia, Kartagener syndrome
<i>Dyx1c1</i>	Primary ciliary dyskinesia, Kartagener syndrome
<i>Cc2d2a</i>	Joubert Syndrome; COACH syndrome
<i>Cep290</i>	Joubert Syndrome 5; Senior-loken Syndrome6, Bardet Biedl Syndrome Leber Congenital Amaurosis 10, Meckel Syndrome-4
<i>Jbts17</i>	Joubert Syndrome
<i>Kif7</i>	Joubert Syndrome, Acrocallosal Syndrome, Hydrolethalus Syndrome
<i>Tmem67</i>	Joubert Syndrome; Nephronophthisis-11, Meckel Syndrome, COACH syndrome, Bardet Biedl Syndrome
<i>Anks6</i>	Polycystic kidney disease; nephronophthisis 16
<i>Bicc1</i>	Polycystic kidney disease; cystic renal dysplasia
<i>Nek8</i>	Polycystic kidney disease; nephronophthisis 9
<i>Pkd1</i>	Polycystic kidney disease
<i>PkdIII</i>	Kidney disease, diaphanospondylodysostosis
<i>Dync2h1</i>	Short rib polydactyly; asphyxiating thoracic dystrophy
<i>Ift140</i>	Mainzer-Saldino Syndrome, short rib thoracic dysplasia
Other Human Diseases	
<i>Ap2b1</i>	Ataxia telangiectasia, cerebellar degeneration
<i>Cfc1</i>	Conotruncal heart defects associated with heterotaxia
<i>Cxcr4</i>	WHIM Syndrome
<i>Dnm2</i>	Charcot-Marie-Tooth Syndrome, Centronuclear myopathy
<i>Frem2</i>	Fraser's Syndrome, Cryptophthalmos syndrome
<i>Lrp1</i>	Alzheimer Disease, Schizophrenia
<i>Lrp2</i>	Donnai-Barrow Syndrome; Facio-oculo-acoustico-renal (FOAR)
<i>Ltbp1</i>	Pseudoexfoliation Syndrome
<i>Megf8</i>	Carpenter's syndrome
<i>Pcsk5</i>	VACTERL/VATER syndrome
<i>Prdm1</i>	Devic disease, B cell lymphoma
<i>Smad6</i>	Aortic valve disease
<i>Smarca4</i>	Rhabdoid tumor predisposition syndrome type 2 (RTPS2), Coffin-Siris syndrome
<i>Snx17</i>	Cerebral cavernous malformation; cavernous malformation
<i>Sufu</i>	Medulloblastoma, basal cell nevus syndrome

Coordination of mitophagy and mitochondrial biogenesis during ageing in *C. elegans*

Konstantinos Palikaras^{1,2}, Eirini Lionaki¹ & Nektarios Tavernarakis^{1,3}

Impaired mitochondrial maintenance in disparate cell types is a shared hallmark of many human pathologies and ageing^{1–8}. How mitochondrial biogenesis coordinates with the removal of damaged or superfluous mitochondria to maintain cellular homeostasis is not well understood. Here we show that mitophagy, a selective type of autophagy targeting mitochondria for degradation, interfaces with mitochondrial biogenesis to regulate mitochondrial content and longevity in *Caenorhabditis elegans*. We find that DCT-1 is a key mediator of mitophagy and longevity assurance under conditions of stress in *C. elegans*. Impairment of mitophagy compromises stress resistance and triggers mitochondrial retrograde signalling through the SKN-1 transcription factor that regulates both mitochondrial biogenesis genes and mitophagy by enhancing DCT-1 expression. Our findings reveal a homeostatic feedback loop that integrates metabolic signals to coordinate the biogenesis and turnover of mitochondria. Uncoupling of these two processes during ageing contributes to overproliferation of damaged mitochondria and decline of cellular function.

A common denominator of numerous pathological conditions and ageing is a deregulation of cellular mitochondrial content^{1–8}. To uncover the cellular and molecular underpinnings of mitochondrial mass homeostasis, we examined the involvement of mitochondria-selective autophagy (mitophagy) in regulating mitochondrial content of *C. elegans* during ageing. Mitochondria gradually accrue with age in wild-type nematodes¹ (Fig. 1a and Extended Data Fig. 1a–d). Depletion of the Atg6/Vps30/Beclin1 homologue BEC-1, the main autophagy regulator in worms⁹, recapitulates the effect of ageing on mitochondrial mass in young adult animals (Fig. 1b). Thus, failure of autophagy impairs the removal of mitochondria and may contribute to progressive accumulation of mitochondria during ageing.

To target mitophagy specifically, we knocked down *dct-1* (DAF-16/FOXO Controlled, germline Tumour affecting-1)^{10,11}, a putative orthologue to the mammalian NIX/BNIP3L and BNIP3 (Nip3-like protein X/Bcl-2 and adenovirus E1B interacting protein; Extended Data Fig. 2a), which act as mitophagy receptors in mammals^{12–15}. DCT-1 is widely expressed throughout development as an integral membrane protein and localized on the outer mitochondrial membrane (Extended Data Fig. 2b–g). Expression of *dct-1* is partly controlled by the FOXO transcription factor DAF-16 and is elevated upon low insulin/IGF-1 signalling^{10,11} (Extended Data Fig. 1e, f). Similarly to BEC-1, DCT-1 deficiency increases intestinal and muscle cell mitochondrial content in young adult animals and distorts mitochondrial network morphology (Fig. 1c, d and Extended Data Figs 1g–k and 3a–c). DCT-1 encompasses a WXXL motif and co-localizes with the nematode Atg8/LC3 autophagosome protein LGG-1 on the outer mitochondrial membrane^{12–15} (Extended Data Figs 2a and 4a). These findings indicate that DCT-1 is the nematode orthologue of NIX/BNIP3L and functions as a key regulator of mitophagy.

To confirm that impairment of mitophagy underlies mitochondrial accumulation in worms, we developed two composite systems for monitoring mitophagy *in vivo*. First, we generated transgenic animals

expressing mitochondria-targeted Rosella (mtRosella) biosensor that combines a fast-maturing pH-insensitive DsRed fused to a pH-sensitive green fluorescent protein (GFP) variant. To validate this tool, we examined animals carrying mtRosella, under normal and mitophagy-inducing conditions (heat stress, carbonyl cyanide *m*-chlorophenylhydrazone (CCCP) and paraquat exposure). All treatments reduce the GFP/DsRed ratio of Rosella fluorescence, indicating stimulation of mitophagy. Knockdown of *lgg-1* gene, a core component of the general autophagic machinery, blocks mitophagy (Extended Data Fig. 4f–h). Additionally, we generated transgenic animals expressing a mitochondria-targeted GFP, together with the autophagosomal marker LGG-1 fused with DsRed in body wall muscle cells, and looked for co-localization under normal and mitophagy-inducing conditions. Mitophagy-inducing stimuli induce the formation of autophagosomes that extensively co-localize with mitochondria

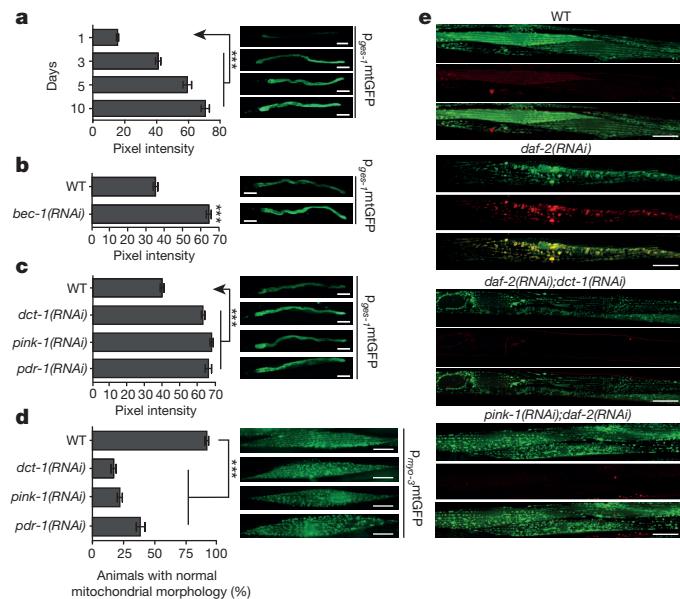
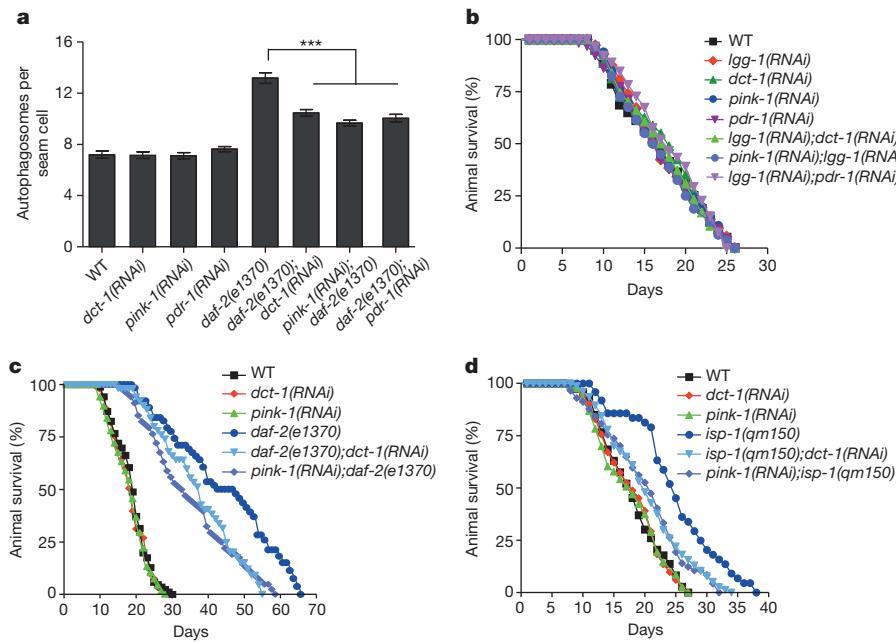


Figure 1 | Selective mitochondrial autophagy regulates mitochondrial content and morphology. Transgenic animals expressing mitochondria-targeted GFP in the intestine were monitored (a) throughout adulthood for mitochondrial content. Indicative images are shown on the right for days 1, 3, 5 and 10 ($n = 90$; *** $P < 0.001$; one-way analysis of variance (ANOVA)). b, BEC-1 deficiency increases mitochondrial mass ($n = 100$; *** $P < 0.001$; unpaired t -test). c, Knockdown of either DCT-1, PINK-1 or PDR-1 increases mitochondrial mass ($n = 100$; *** $P < 0.001$; one-way ANOVA). Scale bar, 100 μ m. d, Depletion of either DCT-1, PINK-1 or PDR-1 alters mitochondrial network morphology in body wall muscle cells ($n = 120$; *** $P < 0.001$; one-way ANOVA). Scale bar, 20 μ m. Anterior, left; posterior, right. e, Induction of mitophagy is indicated by co-localization of GFP and DsRed (for each group of images, mitochondria are shown in green on top, autophagosomes in red below, with a merged image at the bottom). Scale bars, 20 μ m. Images were acquired using $\times 5$ and $\times 40$ objective lenses. Error bars, s.e.m.

¹Institute of Molecular Biology and Biotechnology, Foundation for Research and Technology - Hellas, Nikolaou Plastira 100, Heraklion 70013, Crete, Greece. ²Department of Biology, University of Crete, Heraklion 70013, Crete, Greece. ³Department of Basic Sciences, Faculty of Medicine, University of Crete, Heraklion 71110, Crete, Greece.



(Extended Data Fig. 5a, b). Pronounced induction of mitophagy was also observed in long-lived *daf-2* mutants (Fig. 1e and Extended Data Figs 4f, h and 5b). In all cases, induction of mitophagy is entirely dependent on DCT-1 (Fig. 1e and Extended Data Figs 4f, g and 5a). Additionally, we tested the requirement of the mitochondrial phosphatase and tensin (PTEN)-induced kinase 1 (PINK-1) and the cytosolic E3 ubiquitin ligase Parkin, two critical conserved components of mitophagy¹⁶. Consistent with their previously described role, PINK-1 and PDR-1 (the nematode Parkin homologue) deficiency blocks mitophagy under stress or low insulin/IGF-1 signalling, and destabilizes mitochondrial network in wild-type *C. elegans* cells (Fig. 1c–e

Figure 2 | Mitophagy is required for longevity under conditions of low insulin/IGF-1 signalling or impaired mitochondrial function.

a, Quantification of GFP:LGG-1-positive foci in seam cells of L4 wild-type and *daf-2(e1370)* animals upon knockdown of either *dct-1*, *pink-1* or *pdr-1* ($n = 150$ cells per assay; *** $P < 0.001$; one-way ANOVA). Error bars, s.e.m. **b**, Knockdown of either *dct-1*, *pink-1*, *pdr-1* or *lgg-1* compared with animals subjected to simultaneous knockdown of *lgg-1* and *dct-1* or *lgg-1* and *pink-1* or *lgg-1* and *pdr-1* does not affect the lifespan of otherwise wild-type animals. **c, d**, Knockdown of either *dct-1* or *pink-1* shortens the lifespan of long-lived *daf-2(e1370)*, or *isp-1(qm150)* mutants. Lifespan values are given in Supplementary Table 1; assays were performed at 20 °C.

and Extended Data Figs 1i, j, 4f, g and 5a). Increased mitochondrial mass in mitophagy-depleted animals is not a compensatory protective mechanism against extensive dysfunctional mitochondria or due to enhanced mitochondrial biogenesis, impaired proteasomal activity or altered tissue-specific promoter activity of the reporters (Extended Data Fig. 6a–h). Furthermore, mitochondrial morphology defects upon mitochondrial dynamics impairment are epistatic to morphological aberrations occurring upon mitophagy inhibition, suggesting the latter are secondary events (Extended Data Fig. 3a–c).

Depletion of CED-9, the nematode homologue of cell death inhibitor Bcl-2 and physical interactor of DCT-1 (refs 17, 18), causes mito-

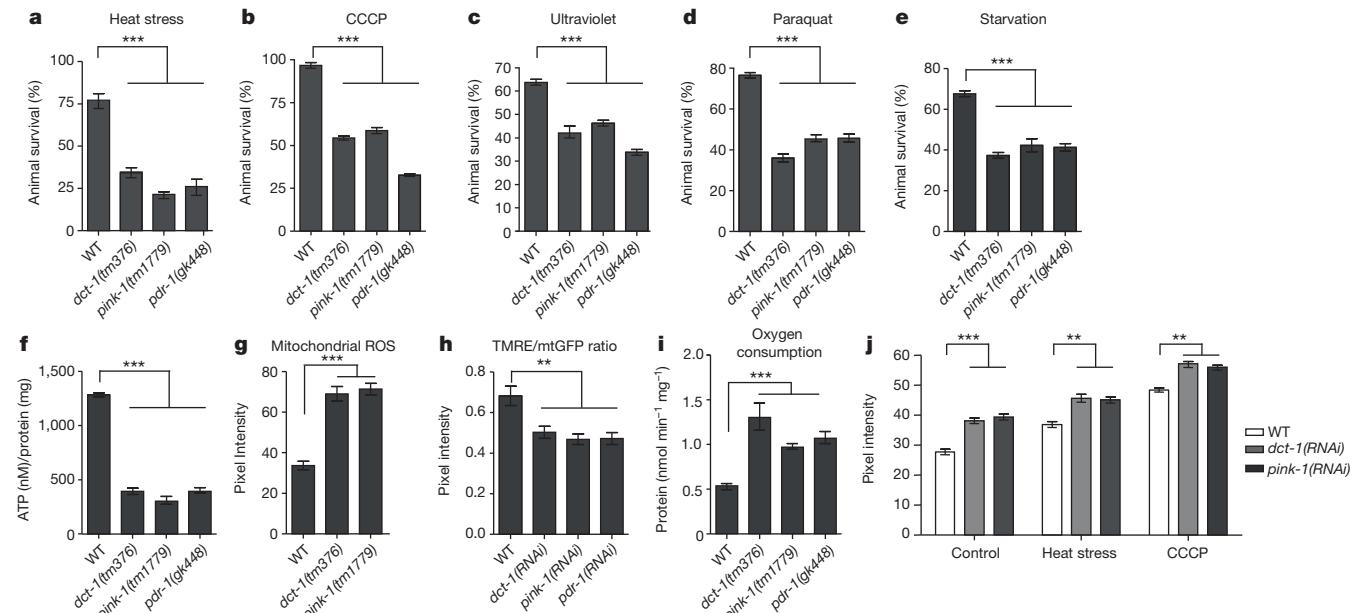


Figure 3 | Mitophagy deficiency compromises stress resistance and impairs mitochondrial function. Mutants *dct-1(tm376)*, *pink-1(tm1779)*, *pdr-1(gk448)* and wild-type animals incubated **a**, at 37 °C for 7 h ($n = 150$; *** $P < 0.001$; one-way ANOVA), **b**, in CCCP for 2 h ($n = 150$; *** $P < 0.001$; one-way ANOVA), or **c**, exposed to 500 J m⁻² of ultraviolet radiation ($n = 150$; *** $P < 0.001$; one-way ANOVA), **d**, treated with 8 mM paraquat ($n = 150$; *** $P < 0.001$; one-way ANOVA). **e**, Survival after 5 days of starvation ($n = 150$; *** $P < 0.001$; one-way ANOVA). **f**, ATP levels are decreased in mitophagy mutants (*** $P < 0.001$; one-way ANOVA).

one-way ANOVA). **g**, Mitophagy-deficient animals display increased mitochondrial ROS formation ($n = 180$; *** $P < 0.001$; one-way ANOVA). **h**, Mitophagy mutants display decreased ratio of dysfunctional mitochondrial to total mitochondrial mass ($n = 170$; ** $P < 0.01$; one-way ANOVA). **i**, Oxygen consumption rates (** $P < 0.01$; one-way ANOVA). **j**, Transgenic animals expressing the Ca^{2+} -reporter upon knockdown of either *dct-1* or *pink-1* incubated at 37 °C for 3 h, with or without the presence of CCCP ($n = 120$; *** $P < 0.001$, ** $P < 0.01$; one-way ANOVA). Error bars, s.e.m.

chondrial mass increase without affecting DCT-1 mitochondrial localization (Extended Data Fig. 7a, b). Interestingly, CED-9 is required for mitophagy upon oxidative and mitochondrial stress (Extended Data Fig. 7c, d). Simultaneous knocking down of *ced-9/dct-1*, *ced-9/pink-1* and *ced-9/pdr-1* does not further increase mitochondrial content, indicating that *ced-9*, *dct-1*, *pink-1* and *pdr-1* act in the same genetic pathway to regulate mitophagy (Extended Data Fig. 7b).

To investigate the contribution of mitophagy to longevity, we knocked down DCT-1, PINK-1 and PDR-1 in wild-type and long-lived mutants. Notably, DCT-1, PINK-1 and PDR-1 impairment functions non-redundantly to abrogate elevated autophagy in long-lived *daf-2* mutants and remarkably shortens their lifespan, whereas it does not alter basal autophagy or lifespan of wild-type animals (Fig. 2a–c, Extended Data Figs 4c–e and 8a–f and Supplementary Table 1). Moreover, *dct-1* or *pink-1* knockdown shortens lifespan extension of mutants with moderated mitochondrial dysfunction (*isp-1* and *clk-1*), or the inherently caloric-restricted *eat-2* mutants (Fig. 2d, Extended Data Fig. 8g, h and Supplementary Table 1). We conclude that mitophagy is a common longevity assurance process, underlying the pro-survival effects of diverse manipulations that extend lifespan.

To elucidate the mechanism by which mitophagy contributes to longevity, we tested mitophagy-depleted animals for stress resistance. The *dct-1*, *pink-1* and *pdr-1* mutants are substantially more sensitive to various stressors (Fig. 3a–e and Extended Data Fig. 9a, b), whereas DTC-1 overexpression confers PINK-1- and PDR-1-dependent protection against stress. Interestingly, the WXXL motif of DTC-1 is required for DCT-1-dependent stress resistance (Extended Data Figs 8l and 9g–k). We hypothesized that diminished stress resistance may originate from the accumulation of dysfunctional mitochondria in mitophagy mutants. Indeed, in mitophagy-deficient animals, although mitochondrial DNA (mtDNA) levels remain relatively unchanged (Extended Data Fig. 9f), we noted decreased ATP levels, elevated mitochondrial reactive oxygen species (ROS) generation, mitochondrial membrane depolarization, increased oxygen consumption and cytoplasmic Ca^{2+} elevation that becomes exacerbated under stress (Fig. 3f–j and Extended Data Fig. 9c–e). These findings, combined with the distorted mitochondrial network morphology noted above, indicate that mitophagy impairment precipitates pronounced mitochondrial dysfunction.

To examine the molecular interaction between DCT-1, PINK-1 and PDR-1, we immunoprecipitated DCT-1 upon mitophagy stimulation, and analysed tryptic peptides by mass spectrometry. We found that DCT-1 is ubiquitinated on lysine 26, a modification that becomes enriched under mitophagy-inducing conditions and depends on PINK-1 (Extended Data Fig. 9l–o). Additionally, DCT-1 co-localizes with PDR-1, indicating that DCT-1 ubiquitination upon mitophagy is PDR-1-dependent (Extended Data Fig. 4b).

Despite the aforementioned defects, mitophagy-depleted animals show normal lifespan in the absence of stress stimuli (Fig. 2b–d and Extended Data Fig. 8a–h). To resolve this apparent paradox, we considered whether accumulation of dysfunctional mitochondria upon mitophagy impairment signals a compensatory retrograde response that induces expression of nuclear genes encoding mitochondrial proteins. Accumulation of damaged mitochondria induces oxidative stress. In mammalian cells, the nuclear factor-erythroid 2-related factor 2 (NRF2/NFE2L2) is activated under oxidative stress conditions and orchestrates the expression of several genes promoting survival^{19,20}. SKN-1 (SKINhead-1), the nematode homologue of NRF2, was recently found to associate with outer mitochondrial membrane proteins²¹. SKN-1 transcriptional activity is pivotal for mitophagy-mediated longevity^{22,23}, and is implicated in the maintenance of mitochondrial homeostasis^{24–26}. Interestingly, SKN-1 becomes specifically activated upon mitophagy inhibition (Fig. 4a and Extended Data Fig. 10a–d). SKN-1 is required for the expression of several mitochondrial biogenesis genes as well as the mitophagy mediator *dct-1* (Fig. 4b and Extended Data Fig. 10e, h). In addition to DAF-16,

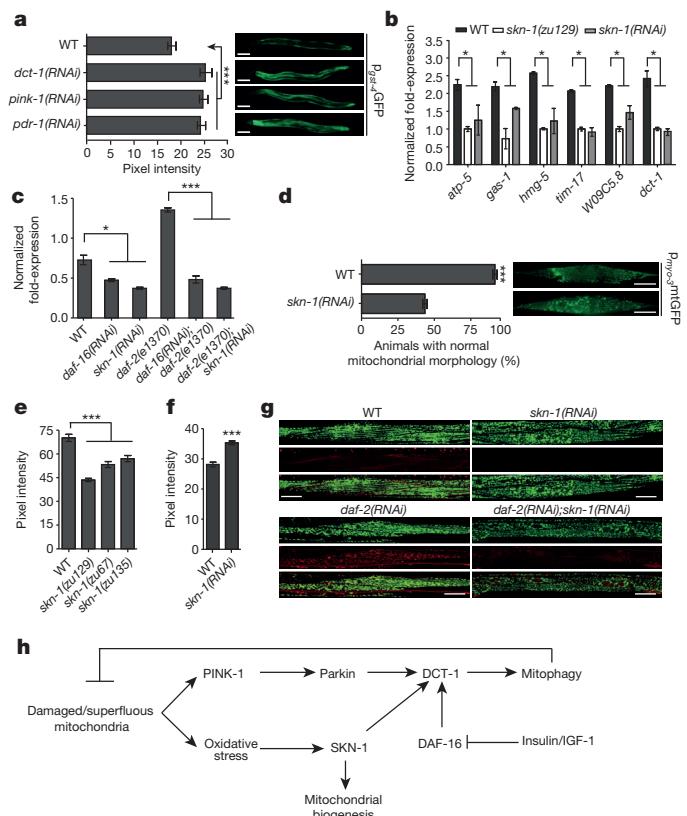


Figure 4 | Mitophagy deficiency engages the mitochondrial retrograde signalling pathway through activation of SKN-1. **a**, Fluorescence intensity of transgenic animals expressing the *pgst-4*GFP transgene subjected to either *dct-1*, *pink-1* and *pdr-1* knockdown ($n = 180$; *** $P < 0.001$; one-way ANOVA). Scale bars, 100 μm . **b**, Messenger RNA (mRNA) levels of *atp-5*, *gas-1*, *hmg-5*, *tim-17*, *W09C5.8* and *dct-1* in wild-type and SKN-1-depleted animals (* $P < 0.01$; one-way ANOVA). **c**, mRNA levels of *dct-1* in animals with the indicated genetic backgrounds (** $P < 0.001$, * $P < 0.01$; one-way ANOVA). **d**, Knockdown of SKN-1 alters mitochondrial network morphology in body wall muscle cells ($n = 120$; *** $P < 0.001$; unpaired t -test). Scale bar, 20 μm . **e**, TMRE staining of wild-type and SKN-1-deficient animals ($n = 150$; *** $P < 0.001$; one-way ANOVA). **f**, Transgenic animals expressing the Ca^{2+} reporter upon knockdown of *skn-1* ($n = 120$; *** $P < 0.001$; unpaired t -test). **g**, Mitophagy is not activated in SKN-1-deficient animals under low insulin signalling conditions (for each group of images mitochondria are shown in green on top, autophagosomes in red below, with a merged image at the bottom). Scale bars, 20 μm . **h**, The contribution of the DCT-1 mitophagy receptor and SKN-1 to mitochondrial physiology and consequently to cellular homeostasis is depicted. DCT-1 is an essential mitophagy mediator that is transcriptionally regulated by both DAF-16 and SKN-1. Under normal conditions, basal-level mitophagy maintains the mitochondrial pool and regulates mitochondrial number, depending on the metabolic state of the cell. Mitophagy is triggered under stress conditions to eliminate damaged mitochondria and promote cell survival. Concomitant activation of SKN-1 enhances mitophagy and promotes mitochondrial biogenesis by stimulating expression of DCT-1 and several mitochondrial genes. Error bars, s.e.m. Images were acquired using $\times 5$ and $\times 40$ objective lenses.

dct-1 expression is also controlled by SKN-1 (Fig. 4c). These findings suggest that SKN-1 may control mitophagy through regulation of *dct-1* transcription. Indeed, *skn-1* knockdown impairs mitochondrial network morphology, causes mitochondrial membrane depolarization and increases cytoplasmic Ca^{2+} concentration (Fig. 4d–f and Extended Data Fig. 10g), similarly to DCT-1 deficiency. Notably, SKN-1 depletion reduces mtDNA content, further highlighting the role of SKN-1 in the maintenance of mitochondrial integrity (Extended Data Fig. 10f). Furthermore, induction of mitophagy under low insulin/IGF-1 signalling or stress is diminished in the absence of SKN-1 (Fig. 4g and Extended Data Fig. 10i, j, l). DAF-16 activity is also

required for mitophagy stimulation. Simultaneous depletion of either DAF-16 and DCT-1, or SKN-1 and DCT-1, does not cause a more pronounced mitophagy defect (Extended Data Fig. 10j), indicating that DAF-16, SKN-1 and DCT-1 are components of a common genetic pathway regulating mitophagy. Knockdown of either *dct-1* or *pink-1* shortens the lifespan of *skn-1* mutant in contrast to wild-type animals (Extended Data Fig. 10k and Supplementary Table 1), indicating that SKN-1 activation compensates the reduced basal mitophagy to maintain normal lifespan. Treatment of mitophagy-deficient animals with the ROS scavengers *N*-acetyl-L-cysteine (NAC) and butylated hydroxyanisole (BHA)^{27,28} does not change their lifespan, ruling out the involvement of mitohormetic feedback signalling (Extended Data Fig. 8i,j and Supplementary Table 1). Surprisingly, supplementation of ethylene glycol tetraacetic acid (EGTA), a calcium-specific chelating agent, shortens the lifespan of mitophagy-deficient worms without influencing lifespan of wild-type animals (Extended Data Fig. 8k). Furthermore, SKN-1 activation is diminished in mitophagy-depleted animals upon EGTA treatment (Extended Data Fig. 10n). Elevated cytoplasmic calcium induces mitochondrial biogenesis through calcium/calmodulin-dependent protein kinase II (CaMKII), p38 mitogen activated protein kinase (MAPK), and PGC-1α^{29,30}. Indeed, UNC-43 (the nematode homologue of CaMKII) regulates the transcriptional activity of SKN-1 upon mitophagy inhibition, where cytoplasmic calcium levels are increased (Extended Data Fig. 10m, o). Our observations suggest that cytoplasmic calcium elevation ensuing upon mitophagy impairment mediates SKN-1 activation, which overcomes the cost of basal mitophagy reduction on longevity.

Overall, our analysis uncovers a key regulatory mechanism that interfaces mitochondrial biogenesis with mitochondrial clearance through selective autophagy, to preserve energy homeostasis in *C. elegans*. A central node of the pathway is the mitophagy mediator DCT-1/NIX, which integrates insulin/IGF-1 signals via DAF-16 and SKN-1, in addition to other stress signals, relayed through SKN-1, to control mitophagy (Fig. 4h). Accumulation of damaged or superfluous mitochondria causes oxidative stress, initiating a bipartite retrograde response that is mediated by SKN-1 and involves the coordinated induction of both mitochondrial biogenesis and mitophagy genes. This closed feedback loop preserves mitochondrial quality by neutralizing damaged, and generating fresh mitochondria (Fig. 4h). Moreover, coordination of biogenesis and turnover of mitochondria allows cells to adjust their mitochondrial content in response to physiological demands, stress and other intracellular or environmental stimuli. In this context, age-related decline of mitophagy both hinders removal of damaged mitochondria and impairs mitochondrial biogenesis to instigate progressive mitochondrial accumulation. The tight evolutionary conservation and ubiquitous expression of the regulatory factors involved in this closely orchestrated response suggest that similar pathways might uphold mitochondrial homeostasis across different taxa.

Online Content Methods, along with any additional Extended Data display items and Source Data, are available in the online version of the paper; references unique to these sections appear only in the online paper.

Received 8 November 2013; accepted 9 February 2015.

Published online 20 April 2015.

1. Artal-Sanz, M. & Tavernarakis, N. Prohibitin couples diapause signalling to mitochondrial metabolism during ageing in *C. elegans*. *Nature* **461**, 793–797 (2009).
2. Campbell, G.R., Ohno, N., Turnbull, D.M. & Mahad, D.J. Mitochondrial changes within axons in multiple sclerosis: an update. *Curr. Opin. Neurol.* **25**, 221–230 (2012).
3. Kaeberlein, M. Lessons on longevity from budding yeast. *Nature* **464**, 513–519 (2010).
4. Malpass, K. Neurodegenerative disease: defective mitochondrial dynamics in the hot seat—a therapeutic target common to many neurological disorders? *Nature Rev. Neurol.* **9**, 417 (2013).
5. Preston, C.C. et al. Aging-induced alterations in gene transcripts and functional activity of mitochondrial oxidative phosphorylation complexes in the heart. *Mech. Ageing Dev.* **129**, 304–312 (2008).

6. Vafai, S.B. & Mootha, V.K. Mitochondrial disorders as windows into an ancient organelle. *Nature* **491**, 374–383 (2012).
7. Wallace, D.C. Diseases of the mitochondrial-DNA. *Annu. Rev. Biochem.* **61**, 1175–1212 (1992).
8. Wredenberg, A. et al. Increased mitochondrial mass in mitochondrial myopathy mice. *Proc. Natl. Acad. Sci. USA* **99**, 15066–15071 (2002).
9. Melendez, A. et al. Autophagy genes are essential for dauer development and lifespan extension in *C. elegans*. *Science* **301**, 1387–1391 (2003).
10. Oh, S.W. et al. Identification of direct DAF-16 targets controlling longevity, metabolism and diapause by chromatin immunoprecipitation. *Nature Genet.* **38**, 251–257 (2006).
11. Pinkston-Gosse, J. & Kenyon, C. DAF-16/FOXO targets genes that regulate tumor growth in *Caenorhabditis elegans*. *Nature Genet.* **39**, 1403–1409 (2007).
12. Sandoval, H. et al. Essential role for Nix in autophagic maturation of erythroid cells. *Nature* **454**, 232–235 (2008).
13. Schweers, R. et al. NIX is required for programmed mitochondrial clearance during reticulocytosis maturation. *Proc. Natl. Acad. Sci. USA* **104**, 19500–19505 (2007).
14. Zhang, J. et al. A short linear motif in BNIP3L (NIX) mediates mitochondrial clearance in reticulocytes. *Autophagy* **8**, 1325–1332 (2012).
15. Zhu, Y. et al. Modulation of serines 17 and 24 in the LC3-interacting region of BNIP3 determines pro-survival mitophagy versus apoptosis. *J. Biol. Chem.* **288**, 1099–1113 (2013).
16. Pickrell, A.M. & Youle, R.J. The roles of PINK1, Parkin, and mitochondrial fidelity in Parkinson's disease. *Neuron* **85**, 257–273 (2015).
17. Xue, D. & Horvitz, H.R. *Caenorhabditis elegans* CED-9 protein is a bifunctional cell-death inhibitor. *Nature* **390**, 305–308 (1997).
18. Yasuda, M., D'Sa-Eipper, C., Gong, X.L. & Chinnadurai, G. Regulation of apoptosis by a *Caenorhabditis elegans* BNIP3 homolog. *Oncogene* **17**, 2525–2530 (1998).
19. Ma, Q. Role of nrf2 in oxidative stress and toxicity. *Annu. Rev. Pharmacol. Toxicol.* **53**, 401–426 (2013).
20. Yun, J. & Finkel, T. Mitohormesis. *Cell Metab.* **19**, 757–766 (2014).
21. Paek, J. et al. Mitochondrial SKN-1/Nrf mediates a conserved starvation response. *Cell Metab.* **16**, 526–537 (2012).
22. Schmeisser, S. et al. Mitochondrial hormesis links low-dose arsenite exposure to lifespan extension. *Aging Cell* **12**, 508–517 (2013).
23. Zarse, K. et al. Impaired insulin/IGF1 signaling extends life span by promoting mitochondrial L-proline catabolism to induce a transient ROS signal. *Cell Metab.* **15**, 451–465 (2012).
24. Ghose, P., Park, E.C., Tabakin, A., Salazar-Vasquez, N. & Rongo, C. Anoxia-reoxygenation regulates mitochondrial dynamics through the hypoxia response pathway, SKN-1/Nrf, and stomatin-like protein STL-1/SLP-2. *PLoS Genet.* **9**, e1004063 (2013).
25. Staab, T.A., Evgrafov, O., Knowles, J.A. & Sieburth, D. Regulation of synaptic nlg-1/neuroligin abundance by the skn-1/Nrf stress response pathway protects against oxidative stress. *PLoS Genet.* **10**, e1004100 (2014).
26. Weimer, S. et al. D-glucosamine supplementation extends life span of nematodes and of ageing mice. *Nature Commun.* **5**, 3563 (2014).
27. Schmeisser, K. et al. Role of sirtuins in lifespan regulation is linked to methylation of nicotinamide. *Nature Chem. Biol.* **9**, 693–700 (2013).
28. Schulz, T.J. et al. Glucose restriction extends *Caenorhabditis elegans* life span by inducing mitochondrial respiration and increasing oxidative stress. *Cell Metab.* **6**, 280–293 (2007).
29. Ventura-Clapier, R., Garnier, A. & Veksler, V. Transcriptional control of mitochondrial biogenesis: the central role of PGC-1α. *Cardiovasc. Res.* **79**, 208–217 (2008).
30. Wright, D.C., Geiger, P.C., Han, D.H., Jones, T.E. & Holloszy, J.O. Calcium induces increases in peroxisome proliferator-activated receptor gamma coactivator-1α and mitochondrial biogenesis by a pathway leading to p38 mitogen-activated protein kinase activation. *J. Biol. Chem.* **282**, 18793–18799 (2007).

Supplementary Information is available in the online version of the paper.

Acknowledgements We thank A. Pasparaki for technical support with experiments, N. Kourtis for the GCMP2.0-expressing *C. elegans* strain, N. Charmilas for the neuronal CTS-1::mCherry-expressing *C. elegans* strain and the *unc-43* RNA interference (RNAi) plasmid, and K. Kounakis for the *rpn-6* RNAi plasmid. We thank B.P. Braeckman for the Hyper-expressing *C. elegans* strain and D. Sieburth for the INVOM::RFP-expressing *C. elegans* strain. We thank R. Devenish for providing the pAS1NB-CS-Rosella plasmid. Mass spectrometry analysis was performed at the Institute of Molecular Biology and Biotechnology Proteomics Facility. We are grateful to M. Aivaliotis for the characterization and relative quantification of DCT-1 tryptic peptides by nanoflow liquid chromatography with tandem mass spectrometry. Some nematode strains used in this work were provided by the *Caenorhabditis Genetics Center*, which is funded by the National Center for Research Resources of the National Institutes of Health, and S. Mitani (National Bioresource Project) in Japan. We thank A. Fire for plasmid vectors. This work was funded by grants from the European Research Council, the European Commission 7th Framework Programme and the Greek General Secretariat for Research and Technology.

Author Contributions K.P., E.L. and N.T. designed and performed experiments. K.P. and N.T. analysed data and wrote the manuscript.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to N.T. (tavernarakis@imbb.forth.gr).

METHODS

No statistical methods were used to predetermine sample size.

Strains and genetics. We followed standard procedures for *C. elegans* strain maintenance³¹. Nematode rearing temperature was kept at 20 °C, unless noted otherwise. The following strains were used in this study: N2: wild-type Bristol isolate, CB1370: *daf-2(e1370)III*, CB4876: *clk-1(e2519)III*, MQ887: *isp-1(qm150)IV*, DA465: *eat-2(ad465)II*, CF1038: *daf-16(mu86)I*, *pink-1(tm1779)II*, *dct-1(tm376)X*, VC1024: *pdr-1(gk448)III*, N2;Ex[p_{myo-3}INVOM::RFP], SJ4103: N2;Is[p_{myo-3}mTGF_P], SJ4143: N2;Is[p_{ges-1}mTGF_P], CL2166: N2;Is[p_{gst-4}GFP], CF1553: N2;Is[p_{sod-3}GFP], SJ4005: N2;Is[p_{hsp-4}GFP]V, SJ4058: N2;Is[p_{hsp-60}GFP]V, EU1: *skn-1(zu67)IV*/nT1(IV;V), EU31: *skn-1(zu135)IV*/nT1(IV;V), EU40: *skn-1(zu129)IV*/nT1(IV;V). To investigate the localization of DCT-1, we examined IR1431: N2;Ex001[p_{dct-1}DCT-1::GFP] and IR1607: N2;Is[p_{myo-3}INVOM::RFP]; Ex001[p_{dct-1}DCT-1::GFP] transgenic animals. The following strains were used to monitor mitophagy process: IR1284: N2;Is[p_{myo-3}mTGF_P];Ex011[p_{lgg-1}DsRed::LGG-1], IR1511: N2;Ex001[p_{myo-3}DsRed::LGG-1;p_{dct-1}DCT-1::GFP] and IR1631: N2;Ex003[p_{myo-3}TOMM-20::Rosella]. To examine mitochondrial morphology and mitochondrial mass under mitophagy-deficient conditions, we used the SJ4103: N2;Is[p_{myo-3}mTGF_P], IR1298: *dct-1(tm376)*;Is[p_{myo-3}mTGF_P] and SJ4143: N2;Is[p_{ges-1}mTGF_P], IR1505: N2;Ex001[p_{unc-119}CTS-1::mCherry] strains. To monitor autophagy induction after inhibition of mitophagy, we used the DA2123: N2;Is[p_{lgg-1}GFP::LGG-1], MAH14: *daf-2(e1370)*;Is[p_{lgg-1}GFP::LGG-1] and IR621: N2;Ex002[p_{lgg-1}DsRed::LGG-1] strain. To monitor DCT-1 and PDR-1 interaction *in vivo*, we used the following transgenic animal: IR1739: N2;Ex001[p_{myo-3}PDR-1::DsRed;p_{dct-1}DCT-1::GFP]. To investigate the protective role of global expression of *dct-1* under mitophagy inducing conditions, we examined the following transgenic animals: IR1431: N2;Ex001[p_{dct-1}DCT-1::GFP], IR1487: *pink-1(tm1779)*;Ex001[p_{dct-1}DCT-1::GFP], IR1733: *pdr-1(gk448)*;Ex001[p_{dct-1}DCT-1::GFP], IR1666: *dct-1(tm376)*;Ex002[p_{let-858}DCT-1], IR1668: *dct-1(tm376)*;Ex001[p_{let-858}DCT-1(ΔMER)], IR1667: *dct-1(tm376)*;Ex003[p_{let-858}DCT-1(ΔBH3)], IR1669: *dct-1(tm376)*;Ex003[p_{let-858}DCT-1(ΔWXXL)]. The following strain was used for monitoring cytoplasmic calcium levels: IR1155: N2;Ex001[p_{let-858}GCaMP2.0]. To monitor the levels of ROS, such as H₂O₂, in mitophagy-deficient animals, we used the JR1S1: N2;Is[p_{rpl-1}-HyPer] strain³².

Molecular cloning. To generate the p_{dct-1}DCT-1::GFP reporter construct, we fused a KpnI fragment, containing the coding sequence of *dct-1* and 2 kb upstream of the coding sequence, amplified from *C. elegans* genomic DNA using the primers 5'-GGATCCCATCTTCTGTACCACGCA-3' and 5'-GGTACCGCTTGACGAT-GATGTTGAT-3', at the amino (N) terminus of GFP, in the pPD95.77 plasmid vector. The translational p_{dct-1}DCT-1::GFP fusion construct was co-injected with pRF4 (contains the *rol-6(su1006)* dominant transformation marker) into the gonads of wild-type animals. To generate the p_{let-858}GCaMP2.0 reporter construct, we removed the *mec-17* promoter form p_{mec-17}GCaMP2.0 described previously³³, and inserted a HindIII/SmaI fragment (derived from the L3786 plasmid vector), containing the sequence of the *let-858* promoter upstream of GCaMP2.0. The resulting plasmid construct was co-injected with pRF4 into the gonad of wild-type animals. To generate p_{myo-3}DsRed::LGG-1, we inserted an AgeI/EcoRI fragment (derived from p_{lgg-1}DsRed::LGG-1³⁴), containing the coding sequence of DsRed downstream of the *myo-3* promoter of the pPD96.52 plasmid vector. We then fused an EcoRI fragment (derived from p_{lgg-1}DsRed::LGG-1³⁴) containing the coding sequence of *lgg-1* at the carboxy (C) terminus of DsRed of the p_{myo-3}DsRed. The translational p_{myo-3}DsRed::LGG-1 fusion construct was co-injected with p_{myo-2}GFP transformation marker. To generate p_{myo-3}TOMM-20::Rosella, we inserted an EcoRI fragment (derived from the mtRosella yeast vector^{35,36}) containing the coding sequence of Rosella biosensor in pPD96.52 under *myo-3* promoter. We then inserted a BamHI fragment containing the two first exons of *tomm-20* gene downstream of the *myo-3* promoter and upstream of the Rosella coding sequence. The *tomm-20* fragment was amplified from *C. elegans* genomic DNA using the primers 5'-GGATCCATGTCGGACACAATTCTTGG-3' and 5'-GGATCCCTCGCGAGTCGACGGC-3'. The translational p_{myo-3}TOMM-20::Rosella fusion construct was co-injected with pRF4 into the gonads of wild-type animals. To generate p_{myo-3}PDR-1::DsRed, we inserted an AgeI/BamHI fragment containing the full length *pdr-1* gene without the stop codon, amplified from *C. elegans* genomic DNA using the primers 5'-GTGCGACTCGCAAGCAGGCCGC-3' and 5'-ACCGGTGGATTAAACCAATGGTCC-3', at the N terminus of DsRed, in the pPD96.52 plasmid vector. The translational p_{myo-3}PDR-1::DsRed fusion construct was co-injected with pRF4 into the gonads of wild-type animals. To generate p_{unc-119}CTS-1::mCherry, we inserted a XbaI/XmnI fragment containing the first two exons of *cts-1* gene, amplified from *C. elegans* genomic DNA using the primers 5'-ATGTCGCTCTCTGGAATGGC-3' and 5'-GGTGCCTACCGTCTGTGC-3', at the N terminus of GFP, in the pPD95.77 plasmid vector. We then inserted a HindIII/XbaI fragment, amplified

from *C. elegans* genomic DNA using the primers 5'-CTCTAAACATTCAAGTTT TTAAACCAC-3' and 5'-ATATGCTGTTGAGCTAAAATTTG-3' containing the sequence of the *unc-119* promoter upstream of CTS-1::GFP. *p_{unc-119}CTS-1* was inserted as a KpnI fragment upstream of mCherry in the pPD95.77 plasmid vector. The translational *p_{unc-119}CTS-1*::mCherry fusion construct was co-injected with pRF4 into the gonads of wild-type animals. To generate *p_{let-858}DCT-1* rescue plasmid, we inserted an XbaI/XhoI fragment containing *dct-1* complementary DNA (cDNA) sequence in pPD96.52. The *dct-1* fragment was amplified from *C. elegans* cDNA using the primers 5'-TCTAGAATGCCTCATTTCTTGAGTTGTC-3' and 5'-CTCGAGTTGACGGATGATGTTGATTAGTTTC-3'. We then removed the *myo-3* promoter from pPD96.52 and inserted a HindIII/XmaI fragment (derived from the L3786 plasmid vector), containing the *let-858* promoter, upstream of *dct-1*. The translational *p_{let-858}DCT-1* fusion construct was co-injected with pRF4 into the gonads of *dct-1(tm376)* animals. To generate the *p_{let-858}DCT-1(ΔWXXL)* plasmid that contains a deletion of WXXL motif, we inserted an XbaI/AgeI fragment fused with an AgeI/XhoI fragment of the *dct-1* cDNA sequence in pPD96.52, spanning the sequence that encodes for WXXL motif. The *dct-1(ΔWXXL)* fragment was amplified from *C. elegans* cDNA using the primers 5'-TCTAGAATGCCTCATTTCTTGAGTTGTC-3' and 5'-ACCG GTCGATTCTGACATACCCAGGTG-3', 5'-ACCGGTGACCAAGCGTAC AGTT-3' and 5'-CTCGAGTTGACGGATGATGTTGATTAGTTTC-3'. We then removed the *myo-3* promoter from pPD96.52 and we inserted a HindIII/XmaI fragment (derived from the L3786 plasmid vector), containing the sequence of the *let-858* promoter upstream of DCT-1(ΔWXXL). The translational *p_{let-858}DCT-1(ΔWXXL)* fusion construct was co-injected with pRF4 into the gonads of *dct-1(tm376)* animals. To generate *p_{let-858}DCT-1(ΔMER)* plasmid that contains a deletion of the MER domain, we inserted an XbaI/AgeI fragment fused with an AgeI/XhoI fragment of *dct-1* cDNA sequence in pPD96.52 spanning the sequence that encodes for MER domain. *dct-1(ΔMER)* fragment was amplified from *C. elegans* cDNA using the primers 5'-TCTAGAATGCCTCATTTCTTGAGTTGTC-3' and 5'-ACCGGTGACGGTACTGCACAAACT-3', 5'-ACCGGTGCTC AACATCGCACAACTC-3' and 5'-CTCGAGTTGACGGATGATGTTGATTAGT TTC-3'. We then removed the *myo-3* promoter from pPD96.52 and inserted a HindIII/XmaI fragment (derived from the L3786 plasmid vector), containing the *let-858* promoter upstream of *dct-1(ΔMER)*. The translational *p_{let-858}DCT-1(ΔMER)* fusion construct was co-injected with pRF4 into the gonads of *dct-1(tm376)* animals. To generate *p_{let-858}DCT-1(ΔBH3)* plasmid that contains a deletion of the BH3 domain, we inserted an XbaI/AgeI fragment fused with an AgeI/XhoI fragment of *dct-1* cDNA sequence in pPD96.52, spanning the sequence that encodes for the BH3 domain. *dct-1(ΔBH3)* fragment was amplified from *C. elegans* cDNA using the primers 5'-TCTAGAATGCCTCATTTCTTGAGTTGTC-3' and 5'-ACCGGTGACTTACTGTTCTAACGCTCTC-3', 5'-ACCGGTGAGCAGTCGGCCA-3' and 5'-CTCGAGTTGACGGATGATGTTGATTAGTT TC-3'. We then removed the *myo-3* promoter from pPD96.52 and inserted a HindIII/XmaI fragment (derived from the L3786 plasmid vector), containing the *let-858* promoter upstream of *dct-1(ΔBH3)*. The translational *p_{let-858}DCT-1(ΔBH3)* fusion construct was co-injected with pRF4 into the gonads of *dct-1(tm376)* animals. For engineering the *skn-1*, *dct-1*, *pdr-1*, *ced-9*, *pbs-5*, *rpn-6*, *fzo-1*, *drp-1*, *isp-1*, *clk-1* and *unc-43* RNAi constructs, gene-specific fragments of interest were obtained by PCR amplification directly from *C. elegans* genomic DNA using the following sets of primers 5'-TCAGTTCACCATCGTCAAACA CCTC-3' and 5'-CTGCTGCTCCAAACGACAAAGCTG-3', 5'-TTGCGTCGG CGAGAACAGC-3' and 5'-CTGCTGCTCCAAACGACAAAGCTG-3', 5'-GTGCAACTGCAAGCAGGCCGC-3' and 5'-ACCGGTGGATTAAACCAATGGT CCCATTG-3', 5'-ATGACACGCTGCACGGCGG-3' and 5'-CTTCAGCTGA ACATCATCCGCC-3', 5'-ACCGGTATGTGGGGCGAGACATTG-3' and 5'-GGGCCACGTCATCAACACCCAGCC-3', 5'-GAAACGAGTTCACGAG AAGACACTA-3' and 5'-GGAAGGTCAAGCATCACCTGAA-3', 5'-ATTGCGTGTGCTGCCG-3' and 5'-ATGGCGTTGGGGAGAGTCC-3', 5'-ATG GAAAATCTCATTCTGTCGTC-3' and 5'-CGCTCTTCTCCGGTTGATG-3', 5'-ATGGCTCTCTGCTAGATCTGGAGG-3' and 5'-AGTGAGATCC TGACAAGGGCAG-3', 5'-ATGTTCCGTGTAATAACCCGTG-3' and 5'-TCT CAGCAATCGCAATAGCTC-3', 5'-CCAAGTTAGTGACAATTACGATGTG-3' and 5'-CTCTTGGAGAACCTCTGGCGACA-3' respectively. The PCR-generated fragments were subcloned into the pL4440 plasmid vector. The resulting constructs were transformed into HT115(DE3) *Escherichia coli* bacteria deficient for RNase III. Bacteria carrying an empty vector were used in control experiments. For *skn-1*, *pbs-5* and *rpn-6* RNAi, the effect can be severe, leading to embryonic lethality of the animals.

mRNA quantification. To quantify *dct-1*, *atp-5*, *hmg-5*, *gas-1*, *tim-17*, *W09C5.8*, *pink-1*, *pdr-1*, *ges-1*, *myo-3* and *unc-119* mRNA levels, total RNA was extracted using the TRIzol reagent (Invitrogen). The following sets of primers were used, respectively: 5'-GGCTCCAACCTTACCACTCC-3' and 5'-GCAAATCCTACT

GCTGCTCC-3', 5'-ACTGGTCGAAGCTGCCGAG-3' and 5'-ACTCGGCT GGAACCTCTCCG-3', 5'-CGTCCAAGTGTCTCCAAGTG-3' and 5'-CTT CGCTTCGTCGTGTACTCTTT-3', 5'-CTGGATGTCGAGGAACGTG-3' and 5'-AGCTGACACAAGTCCGATGTC-3', 5'-TGCAGAGAAGTACGAATG AGAACGCAC-3' and 5'-CGAGACGAATCCTTCATCTTG-3', 5'-ATGC TGCCACGTTGGCTCTCC-3' and 5'-ACGGAAGCTGTAGCGGTACAAGA GC-3', 5'-AGCATATCGAAATCGAAATGAGTTAG-3' and 5'-TCGACCGTG GCGAGTTACAAG-3', 5'-AGCCACCGAGCGATTGATGTC-3' and 5'-GTG GCATTTGGGCATCTCTTG-3', 5'-TCGCAAGAGGTATGCTTCACAA G-3' and 5'-TGCTGCTCTGCACTGTATCCC-3', 5'-AGCCAGACGGTGAG GAAGATGC-3' and 5'-TGCCTCAAGGGTCTTACATCTGG-3', 5'-AAC ACAACAATCGATCGCAC-3' and 5'-CACGTGTTCCGTGTCGAGGTC-3'. Results were normalized to genomic DNA using the following primers specific for *ama-1*: 5'-CCTACCTACACTCCAAGTCCATCG-3' and 5'-GGTGAAGCT GGCAGATACGTTG-3' and *pmp-3*: 5'-ATGATAAATCAGCGTCCGAC-3' and 5'-TTGCAACGAGAGCAACTGAAC-3'. For cDNA synthesis, mRNA was reverse transcribed using a iScriptTM cDNA Synthesis Kit (BioRad) and PrimeScriptTM Reverse Transcriptase (Takara). Quantitative PCR was performed in triplicate using a Bio-Rad CFX96 Real-Time PCR system (Bio-Rad).

Ca²⁺ monitoring. For intracellular Ca²⁺ monitoring experiments, transgenic animals expressing the Ca²⁺ reporter GCaMP2.0 globally were examined under a Zeiss AxioImager Z2 epifluorescence microscope. L4 larvae were grown on RNAi plates with or without the presence of 15 μM CCCP. RNAi-treated animals against *dct-1* and *pink-1* were monitored on day 4 of adulthood. For heat stress assay, L4 larvae were grown on RNAi plates. On day 4 of adulthood, worms were monitored after 3 h of incubation at 37 °C. Worms were immobilized with levamisole before mounting on 2% agarose pads for microscopic examination with a Zeiss AxioImager Z2 epifluorescence microscope. Images were acquired under the same exposure. Average pixel intensity values were calculated by sampling images of different animals. We calculated the mean and maximum pixel intensity for each animal in these images using the ImageJ software (<http://rsb.info.nih.gov/ij/>). Mean values were compared using unpaired *t*-tests. For multiple comparisons, we used the one-factor (ANOVA) variance analysis corrected by the post hoc Bonferroni test. For each experiment, at least 100 animals were examined for each strain/condition. Each assay was repeated at least three times. We used the Prism software package (GraphPad Software) for statistical analyses.

Mitochondrial imaging. For MitoTracker staining, 500 μl heat-inactivated OP50 (65 °C, 30 min) were mixed with 100 μl MitoTracker Red FM (catalogue number M-22425; Molecular Probes, Invitrogen) stock solution (100 μM) and seeded on agar plates. Animals were incubated for 2 days with the compound. Stained and washed worms were immobilized with levamisole before mounting on 2% agarose pads for microscopic examination with a Zeiss AxioImager Z2 epifluorescence microscope. For the quantification of mitochondrial ROS, 500 μl heat-inactivated OP50 (65 °C, 30 min) were mixed with 100 μl MitoTracker Red CM-H₂X ROS (catalogue number M-7513; Molecular Probes, Invitrogen) stock solution (100 μM) and seeded on agar plates. Animals were incubated for 2 days with corresponding compound. Stained and washed worms were immobilized with levamisole before mounting on 2% agarose pads for microscopic examination with a Zeiss AxioImager Z2 epifluorescence microscope. For TMRE staining (tetramethylrhodamine, ethyl ester, perchlorate; a dye that accumulates in intact, respiring mitochondria), 1-day-adult animals were grown at 20 °C or 25 °C in the presence of 150 nM TMRE for 24 h. Stained and washed worms were immobilized with levamisole before mounting on 2% agarose pads for microscopic examination with a Zeiss AxioImager Z2 epifluorescence microscope. Images were acquired under the same exposure. Average pixel intensity values were calculated by sampling images of different animals. We calculated the mean and maximum pixel intensity for each animal in these images using the ImageJ software (<http://rsb.info.nih.gov/ij/>). For each experiment, at least 150 animals were examined for each strain/condition. Each assay was repeated at least three times. We used the Prism software package (GraphPad Software) for statistical analyses.

Stress assays. To evaluate thermotolerance (heat stress assay), 4-day-old adult hermaphrodites were placed on pre-warmed (37 °C) NGM plates and incubated at 37 °C. After incubation for 7 h, plates were removed; worms were recovered at 20 °C for 1 h. Then, worms were scored for motility, provoked movement and pharyngeal pumping. Worms failing to display any of these traits were scored as dead. Three distinct populations of 150 adults (for each strain) were scored repeatedly over the assay period. For CCCP stress assays, 4-day-old adult hermaphrodites were incubated for 2 h at 20 °C in M9 buffer containing CCCP at a final concentration of 15 μM. Animals were scored for survival after a 1-h recovery period at 20 °C. Three distinct populations of 150 adults (for each strain) were scored repeatedly over the assay period. For ultraviolet stress assays, 4-day-old adult hermaphrodites were exposed to 500 J m⁻² ultraviolet irradiation using a ultraviolet crosslinker (BIO-LINK-BLX-E365, Vilber Lourmat). Three days after

exposure, animals were scored for survival. Three distinct populations of 150 adults (for each strain) were scored repeatedly over the assay period. For oxidative stress assays, 7-day-old adult hermaphrodites were exposed to paraquat at a final concentration of 8 mM. Animals were scored for survival after 2 days at 20 °C. Three distinct populations of 150 adults (for each strain) were scored repeatedly over the assay period. For the starvation assay, worm eggs were collected upon hypochlorite treatment of gravid adults, resuspended in M9 buffer and incubated at 20 °C. Newly hatched worms were arrested in L1 larval stage owing to lack of food. Survival under starvation was assessed every 2 days by taking aliquots and scoring live animals. Three distinct populations (for each strain) were scored repeatedly over the assay period. We used the Prism software package (GraphPad Software) for statistical analyses.

ATP measurements. To determine ATP content, 50 4-day-old adult hermaphrodites were collected in 50 ml of M9 buffer and frozen at -80 °C. Frozen worms were immersed in boiling water for 15 min, cooled and centrifuged to pellet insoluble debris. The supernatant was moved to a fresh tube and diluted tenfold before measurement. ATP content was determined by using the Roche ATP bioluminescent assay kit HSII (Roche Applied Science) and a TD-20/20 luminometer (Turner Designs). ATP levels were normalized to total protein content. We used the Prism software package (GraphPad Software) for statistical analyses.

Mitochondrial DNA quantification. mtDNA was quantified using quantitative real time PCR as described previously³⁸. We used the primers 5'-GTTTATGCT GCTGTAGCGTG-3' and 5'-CTGTTAAAGCAAGTGGACGAG-3' (Mito1 set) for mtDNA. The results were normalized to genomic DNA using the following primers specific for *ama-1*: 5'-TGGAACTCTGGAGTCACACC-3' and 5'-CATC CTCCCTCATTGAACGG-3'. Quantitative PCR was performed using a Bio-Rad CFX96 Real-Time PCR system, and was repeated three times.

Oxygen consumption rate measurements. Oxygen consumption rates were measured as previously described^{37,38} using a Clark-type electrode with some minor modifications (Hansatech Instruments). Young adult worms were washed and collected in M9 buffer. Approximately 100 μl of slurry pellet of worms were delivered into the chamber in 1 ml of M9 buffer. The chamber was kept at 20 °C, and measurements were done for 10–15 min, depending on the oxygen consumption rate. The slope of the straight portion of the plot was used to derive the oxygen consumption rate. Worms were recovered after respiration measurements and collected for protein quantification. Rates were normalized to protein content. We performed three independent measurements per strain. We used the Prism software package (GraphPad Software) for statistical analysis.

Western blotting. Protein samples were analysed by 8% Tricine-SDS-polyacrylamide gel electrophoresis (SDS-PAGE), transferred on nitrocellulose membrane and blotted against various antibodies. The antibodies used in the study were anti-GFP (rabbit polyclonal, Minotech Biotechnology), anti-HSP-60 (mouse monoclonal-Developmental Studies Hybridoma Bank) and anti-β-tubulin (mouse monoclonal-Abcam). For total worm protein extraction, protein samples were produced by directly boiling worms in 2× Laemmli sample buffer with β-mercaptoethanol supplemented with 5 mM PMSF and complete mini proteinase inhibitor cocktail (Roche).

Mitochondria isolation. Mixed populations of worms were incubated in 10 mM DTT in M9 for 30 min at 4 °C with rotation. DTT was washed off with three successive centrifugations and dilutions in M9. Worm bodies were homogenized by hand with a 3 ml Potter-Elvehjem homogenizer with PTFE pestle and glass tube (Sigma-Aldrich) in isolation buffer (50 mM Tris-HCl pH 7.4, 210 mM mannitol, 70 mM sucrose, 0.1 mM EDTA, 2 mM PMSF, and complete mini protease inhibitor cocktail (Roche), the last two being added just before use). The progress of homogenization was monitored by visual inspection under a stereoscope. The suspension was centrifuged at 200g for 1 min and the supernatant fraction kept separately while the pellet was resuspended in isolation buffer and homogenized as previously. The homogenate was centrifuged again at 200g for 1 min. Supernatants were combined and centrifuged once more at low speed (200g) for 1 min and then at high speed (12,000g). The pellet was resuspended in isolation buffer and centrifuged at 200g for 1 min. An aliquot of the supernatant was kept separately for protein concentration determination. High-speed centrifugation followed, and mitochondrial pellets were used directly for proteinase K treatment and carbonate extraction and/or SDS-PAGE.

Proteinase K treatment. Fifty micrograms of crude mitochondria from mixed worm population were digested with 100 μg ml⁻¹ proteinase K (Sigma-Aldrich) in isolation buffer minus protease inhibitors for 30 min. The reaction was inhibited with 5 mM PMSF for 10 min on ice. Proteinase K-shaved mitochondria were analysed by SDS-PAGE and western blotting.

Sodium carbonate extraction. Fifty micrograms of crude mitochondria from mixed worm population were diluted in 0.1 M sodium carbonate (pH 11.5), and incubated on ice for 30 min. The suspension was centrifuged for 30 min at 55,000g

in a Beckman Optima TLX-120 ultracentrifuge. Soluble and membrane fractions were analysed by western blotting.

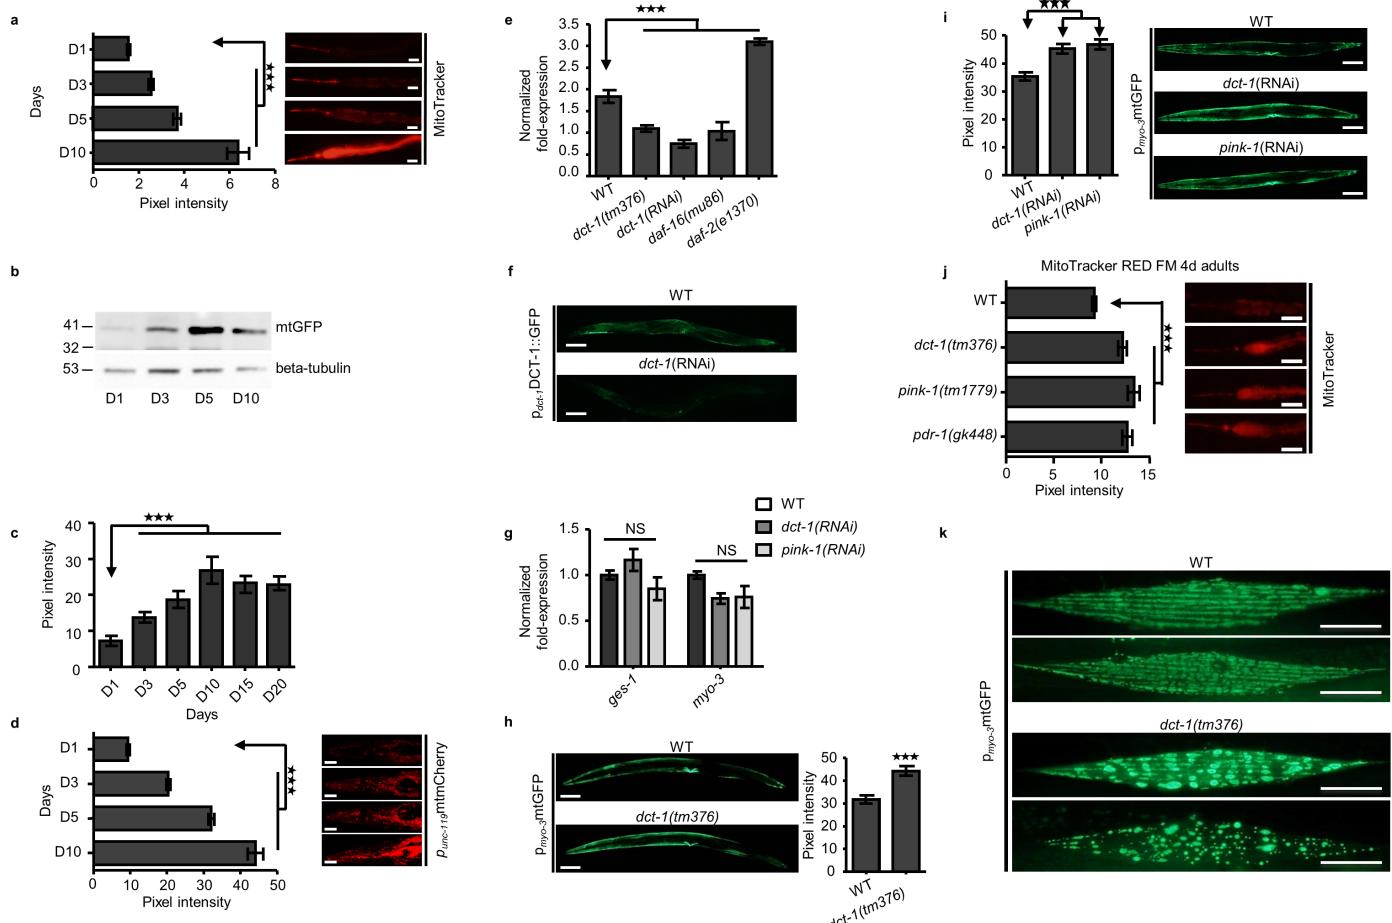
Immunoprecipitation. For immunoprecipitation, mixed populations of worms were treated for 3 h with 10 µM MG132 (Sigma-Aldrich) in M9 buffer and subsequently with 10 µM MG132 and 60 mM paraquat (methyl viologen dichloride, Sigma Aldrich) in M9 buffer. Worm lysis was performed by hand with pellet pestle (Sigma-Aldrich) in lysis buffer (50 mM HEPES-KOH, pH 7.5, 150 mM NaCl, 1 mM EDTA, 1 mM EGTA, 10% glycerol, 1% Triton-X100, 1% NP-40, 0.1% SDS) supplemented with 1 mM DTT, 5 mM PMSF (Sigma-Aldrich), protease and phosphatase inhibitor cocktails (cComplete mini protease inhibitors cocktail tablets-ROCHE, phosphatase inhibitors cocktail-Cell Signaling), 20 µM NEM (Sigma-Aldrich) and 10 µM MG132 just before use. The lysate was incubated on ice for 30 min and then centrifuged for 15 min at 16,000g at 4 °C. Supernatant was diluted three times in dilution buffer (50 mM HEPES-KOH, pH 7.5, 150 mM NaCl, 1 mM EDTA, 1 mM EGTA, 10% glycerol) supplemented with 5 mM PMSF (Sigma-Aldrich), protease and phosphatase inhibitor cocktails (cComplete mini protease inhibitors cocktail tablets (Roche), phosphatase inhibitors cocktail (Cell Signaling), 20 µM NEM (Sigma-Aldrich) and 10 µM MG132 just before use. An equal amount of protein from each sample was incubated for 4 h at 4 °C with Sepharose-Protein G beads (Amersham-GE Healthcare) pre-loaded with rabbit anti-GFP antibody (Minotech Biotechnology), cross-linked onto protein G with BS³ cross-linker (Thermo Scientific). The immunoprecipitated proteins were digested by trypsin for mass spectrometry analysis.

Sample preparation and mass spectrometry analysis. Sepharose beads carrying the immunoprecipitated material were washed twice in 50 mM ammonium bicarbonate buffer. Subsequently, 1 µg of trypsin (sequencing grade (Roche Applied Science)) was added per sample and digestion was performed on beads for 14 h at 37 °C. Tryptic peptides were dried and analysed by nanoflow liquid chromatography-electrospray ionization with tandem mass spectrometry (nLC-ESI-MS/MS) on an EASY-nLC system (Thermo Scientific, software version 2.7.6 #1) coupled with an LTQ-Orbitrap XL ETD (Thermo Scientific) through a nES ion source (Thermo Scientific). Data were acquired with Xcalibur software (LTQ Tune 2.5.5 sp1, Thermo Scientific). The nLC-MS/MS raw data were loaded in Proteome Discoverer 1.3.0.339 (Thermo Scientific) and run using both Mascot 2.3.02 (Matrix Science) and Sequest (Thermo Scientific) search algorithms against the *C. elegans* theoretical proteome (last modified 11 June 2014, version 108). Mass spectrometry was performed in-house in the Proteomics Facility at Institute of Molecular Biology and Biotechnology, Foundation for Research and Technology - Hellas (<http://profi.imbb.forth.gr>).

Lifespan assays. Lifespan assays were performed at 20 °C unless noted otherwise. Synchronous animal populations were generated by hypochlorite treatment of gravid adults to obtain tightly synchronized embryos that were allowed to develop into adulthood under appropriate, defined conditions. The antioxidant NAC was administered at a final concentration of 10 mM. NAC was added into NGM media from a concentrated stock solution (500 mM) before plate pouring. The

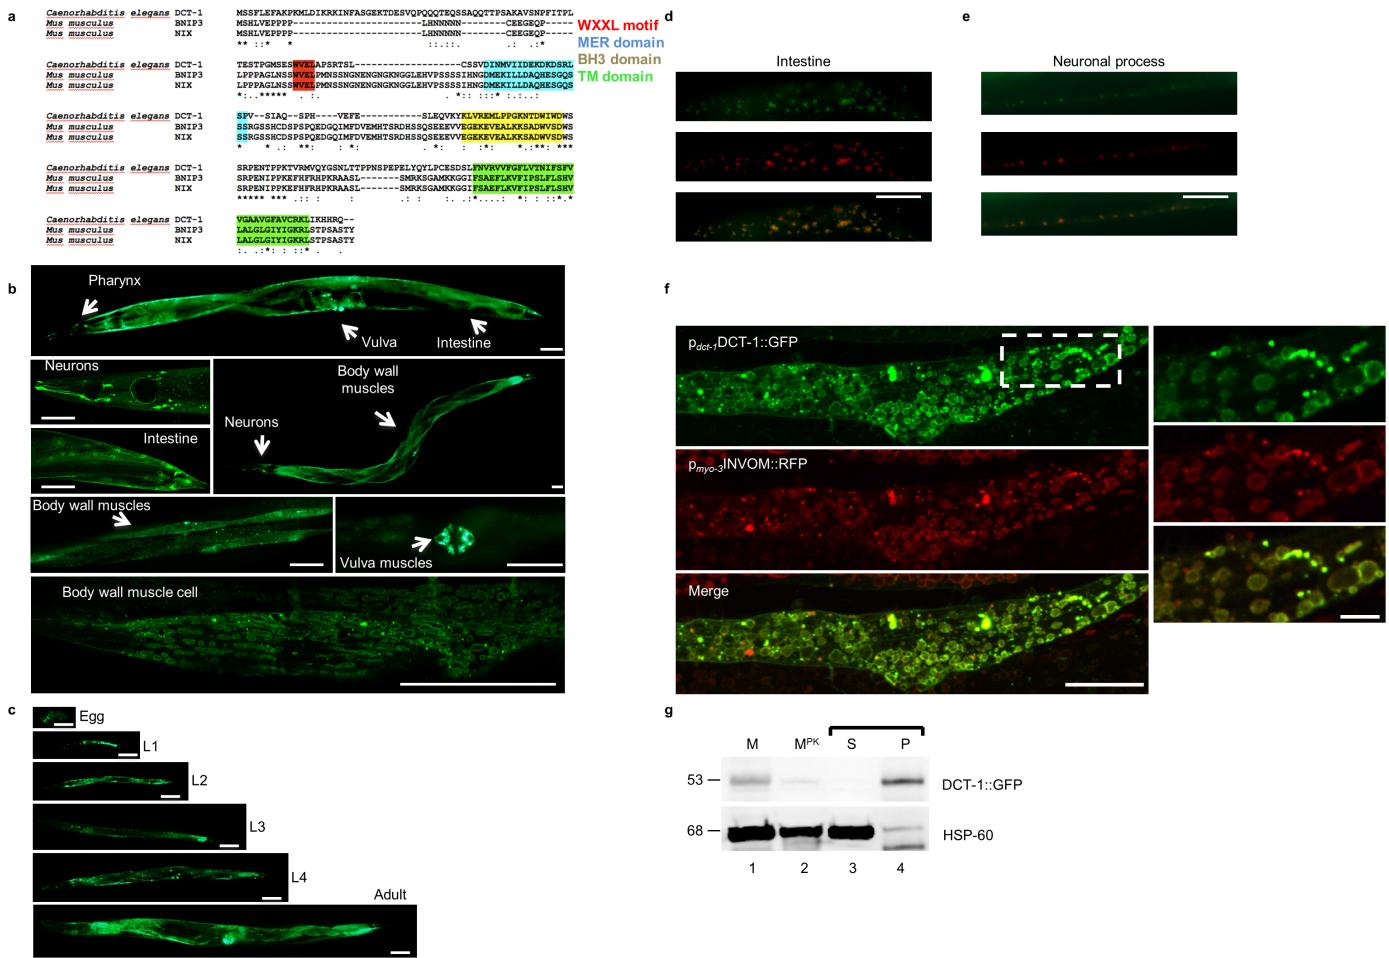
antioxidant butylated hydroxyanisole (BHA) was administered at a final concentration of 25 µM. BHA was added into NGM media from a concentrated stock solution (50 mM) before plate pouring. Ethidium bromide (EtBr) was administered at a final concentration of 5 µg ml⁻¹. EtBr was added into NGM media from a concentrated stock solution (10 mg ml⁻¹) before plate pouring. Ethylene glycol tetraacetic acid (EGTA) was administered at a final concentration of 10 mM. EGTA concentration was prepared by dilutions in 150 µl of sterilized water, from a concentrated stock solution (0.5 M), and applied to the top of the agar medium. Plates were then gently swirled to allow the drug to spread to the entire NGM surface. Identical drug-free water solutions were used for the control plates. For RNAi lifespan experiments, worms were placed on NGM plates containing 2 mM IPTG and seeded with HT115(DE3) bacteria transformed with either the pL4440 vector or the test RNAi construct. Progeny were grown at 20 °C unless noted otherwise, through the L4 larval stage and then transferred to fresh plates in groups of 10–20 worms per plate for a total of 100–150 individuals per experiment. The day of egg collection and initiation of RNAi was used as *t* = 0. Animals were transferred to fresh plates every 2–4 days thereafter and examined every day for touch-provoked movement and pharyngeal pumping, until death. Worms that died owing to internally hatched eggs, an extruded gonad or desiccation due to crawling on the edge of the plates were censored and incorporated as such into the data set. Each survival assay was repeated at least twice and figures represent typical assays. Survival curves were created using the product-limit method of Kaplan and Meier. The log-rank (Mantel-Cox) test was used to evaluate differences between survivals and determine *P* values. We used the Prism software package (GraphPad Software) for statistical analysis and to determine lifespan values¹.

31. Brenner, S. The genetics of *Caenorhabditis elegans*. *Genetics* **77**, 71–94 (1974).
32. Back, P. et al. Exploring real-time *in vivo* redox biology of developing and aging *Caenorhabditis elegans*. *Free Radic. Biol. Med.* **52**, 850–859 (2012).
33. Kourtis, N., Nikoletopoulou, V. & Tavernarakis, N. Small heat-shock proteins protect from heat-stroke-associated neurodegeneration. *Nature* **490**, 213–218 (2012).
34. Samara, C., Syntichaki, P. & Tavernarakis, N. Autophagy is required for necrotic cell death in *Caenorhabditis elegans*. *Cell Death Differ.* **15**, 105–112 (2008).
35. Mijaljica, D., Prescott, M. & Devenish, R. J. A fluorescence microscopy assay for monitoring mitophagy in the yeast *Saccharomyces cerevisiae*. *J. Vis. Exp.* <http://dx.doi.org/10.3791/2779> (2011).
36. Rosado, C. J., Mijaljica, D., Hatzinisiriou, I., Prescott, M. & Devenish, R. J. Rosella: a fluorescent pH-biosensor for reporting vacuolar turnover of cytosol and organelles in yeast. *Autophagy* **4**, 205–213 (2008).
37. Braeckman, B. P., Houthooft, K., De Vreese, A. & Vanfleteren, J. R. Assaying metabolic activity in ageing *Caenorhabditis elegans*. *Mech. Ageing Dev.* **123**, 105–119 (2002).
38. Cristina, D., Cary, M., Lunceford, A., Clarke, C. & Kenyon, C. A regulated response to impaired respiration slows behavioral rates and increases lifespan in *Caenorhabditis elegans*. *PLoS Genet.* **5**, e1000450 (2009).



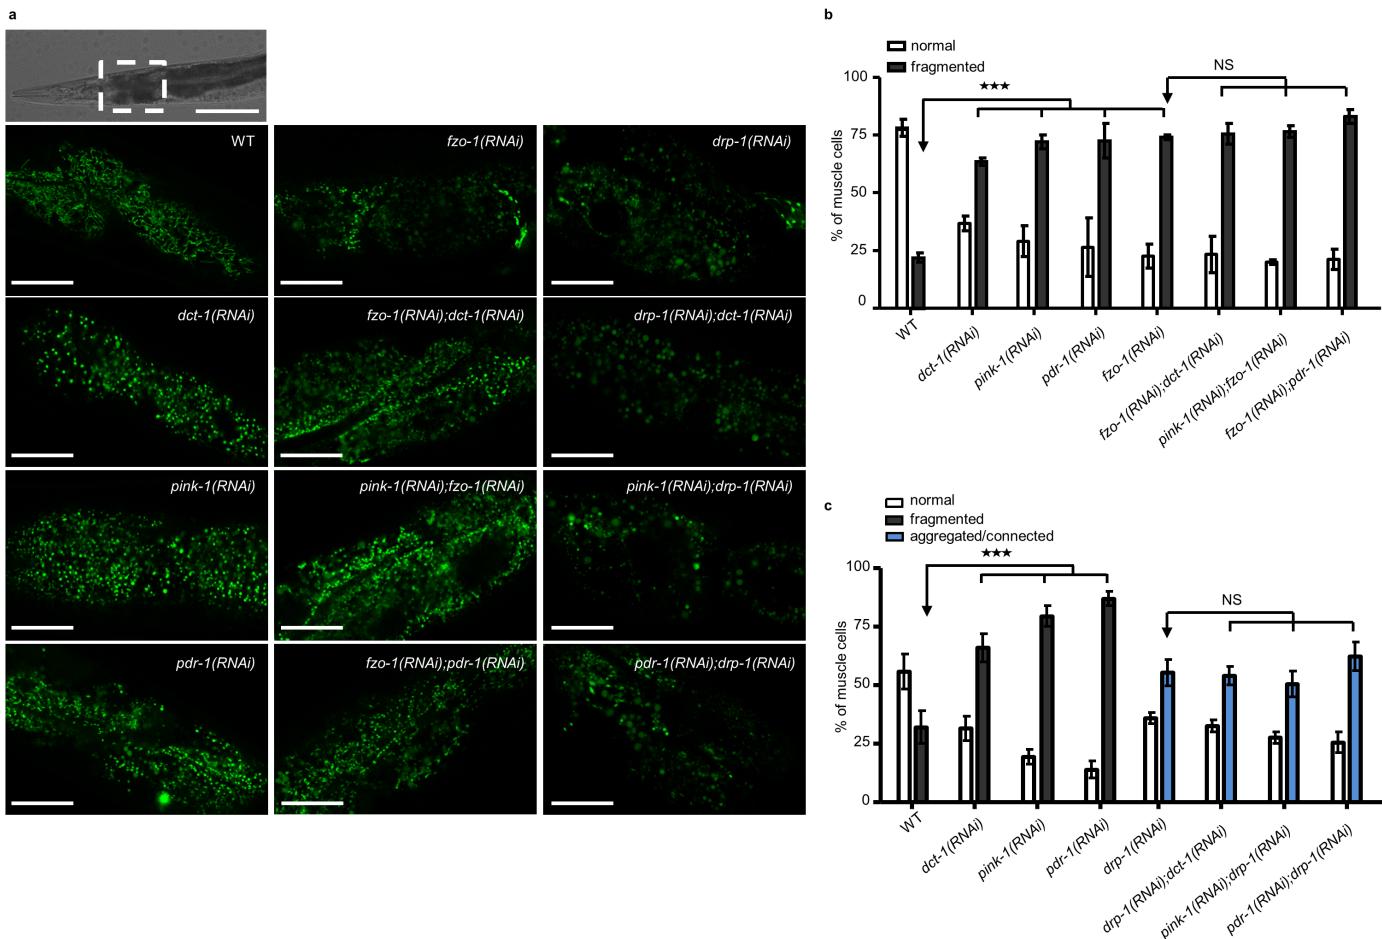
Extended Data Figure 1 | Mitochondrial accumulation during ageing and upon mitophagy depletion. **a**, Stained wild-type animals with MitoTracker Red FM dye were monitored throughout adulthood for mitochondrial content. Indicative images are shown on the right for days 1, 3, 5 and 10 ($n = 90$; *** $P < 0.001$; one-way ANOVA). **b**, Total worm lysates from transgenic animals expressing a mitochondria-tagged GFP in the intestine on day 1, 3, 5 and 10 of adulthood were loaded on Tricine-SDS-PAGE and immunoblotted against the indicated proteins. **c**, Transgenic animals expressing mitochondria-targeted GFP in the intestine were monitored throughout adulthood for mitochondrial content. ($n = 100$; *** $P < 0.001$; one-way ANOVA). **d**, Transgenic animals expressing mitochondria-targeted mCherry in neurons were monitored throughout adulthood for mitochondrial content. Indicative images are shown on the right for days 1, 3, 5 and 10 ($n = 70$; *** $P < 0.001$; one-way ANOVA). **e**, mRNA levels of *dct-1* in wild-type animals, *dct-1(tm376)* mutants, *dct-1(RNAi)* animals, *daf-2(e1370)* and *daf-16(mu86)* mutants. Expression of *dct-1* is transcriptionally regulated by DAF-16 (** $P < 0.001$; one-way ANOVA). **f**, Efficacy of *dct-1* silencing by RNAi. Transgenic animals expressing a full-length *p_{dct-1}DCT-1::GFP* reporter fusion subjected to *dct-1*

RNAi. **g**, mRNA levels in *ges-1* and *myo-3* in wild-type animals, *dct-1(RNAi)* and *pink-1(RNAi)* animals. Expression of *ges-1* and *myo-3* is not changed upon mitophagy inhibition (NS, $P > 0.05$; one-way ANOVA). **h**, Mutant *dct-1(tm376)* animals expressing mitochondria-targeted GFP in body wall muscle cells display increased mitochondrial mass ($n = 100$; *** $P < 0.0001$; unpaired *t*-test). Scale bars, 100 μ m. Images were acquired using a $\times 5$ objective lens. **i**, Knockdown of either DCT-1 or PINK-1 increases mitochondrial mass in worms that express a mitochondria-targeted GFP in body wall muscles ($n = 120$; *** $P < 0.001$; one-way ANOVA). Anterior, left; posterior, right. **j**, Stained wild-type animals and *dct-1(tm376)*, *pink-1(tm1779)* and *pdr-1(gk448)* mutants with MitoTracker Red FM dye were monitored for mitochondrial content ($n = 120$; *** $P < 0.001$; one-way ANOVA). Scale bars, 20 μ m. Images were acquired using a $\times 10$ objective lens. **k**, Mitochondrial network morphology is altered in *dct-1(tm376)* mutants. The mitochondrial network in wild-type animals is well-organized and runs parallel with the myofilament lattice. By contrast, *dct-1(tm376)* mutant animals display fragmented and disorganized mitochondrial network morphology. Scale bar, 20 μ m. Images were acquired using a $\times 40$ objective lens. Error bars, s.e.m.



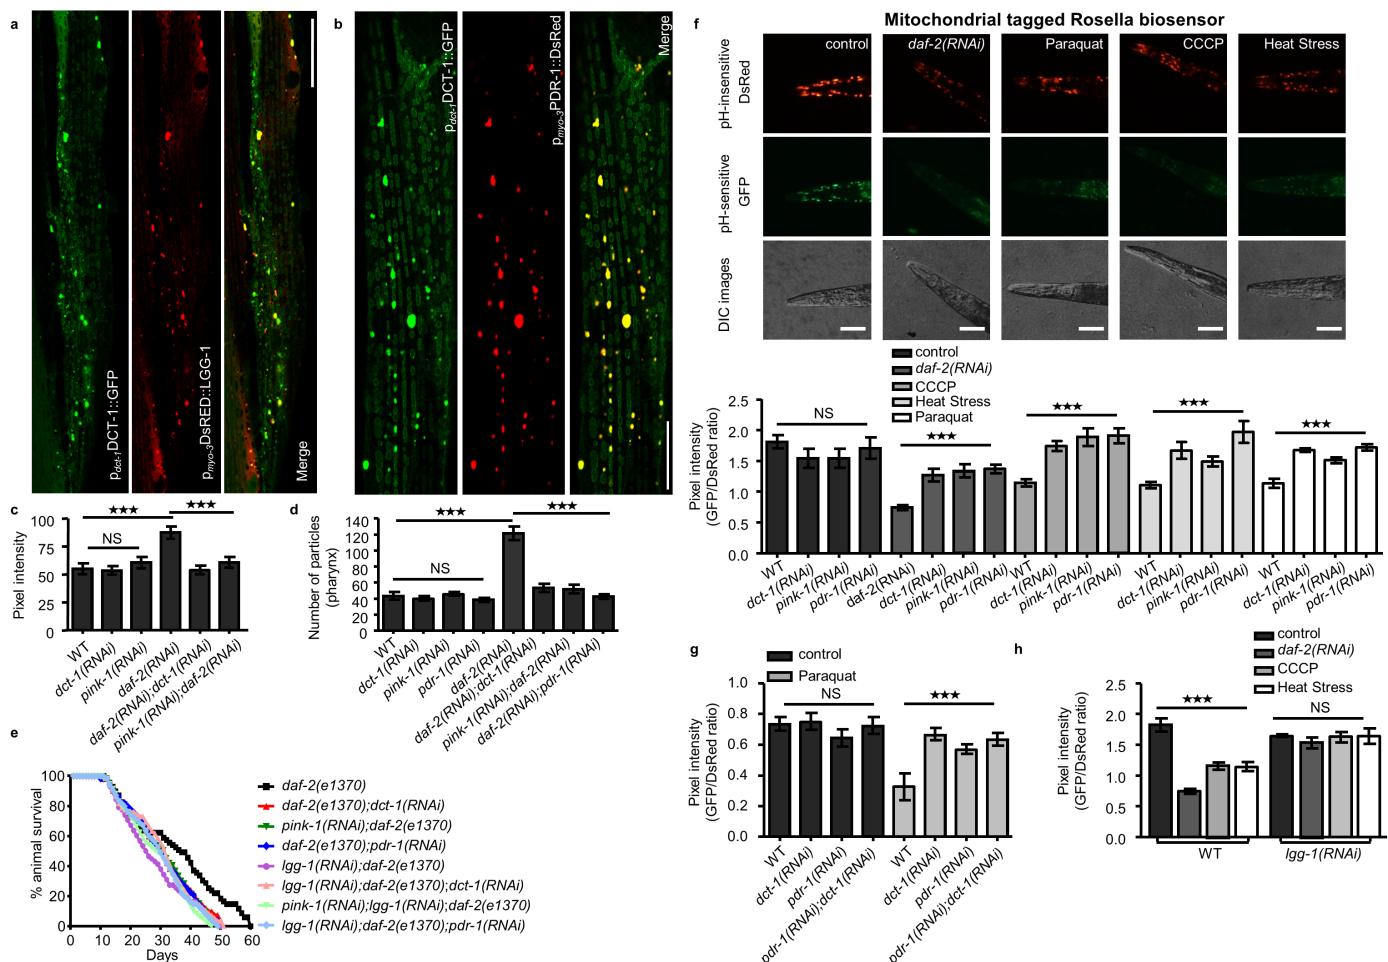
Extended Data Figure 2 | DCT-1 expression and sub-cellular localization.
a, DCT-1 is the homologue of the mammalian BNIP3 and NIX/BNIP3L in *C. elegans*. Alignment of *C. elegans* DCT-1, *M. musculus* BNIP3 and *M. musculus* NIX/BNIP3L proteins. The predicted transmembrane domain of DCT-1, BNIP3 and NIX/BNIP3L is depicted in green. The conserved WXXL motif that is essential for interaction with the autophagosomal protein LC3/Atg8 is shown in red. The conserved MER and BH3 domains are depicted in blue and yellow, respectively. **b**, Spatiotemporal expression and sub-cellular localization of DCT-1. Images of transgenic animals expressing a full-length *p_{dct-1}*:DCT-1::GFP reporter fusion. Expression of *dct-1* occurs in all somatic tissues of adult animals, including neurons, the pharynx, the intestine, body wall muscles and vulva muscles (indicated by the arrows). Scale bars, 50 μm. **c**, Expression is detectable in embryos and remains high during all postembryonic developmental stages, throughout adulthood. Scale bars, 50 μm. Images were acquired using × 10 and × 40 objective lenses. **d**, **e**, DCT-1::GFP co-localizes with mitochondria in different tissues. Transgenic animal expressing a full-length *p_{dct-1}*:DCT-1::GFP reporter fusion localizing in mitochondria. Animals

were stained with Mitotracker Red FM (Invitrogen, Molecular Probes). Scale bars, 20 μm. Images were acquired using a × 100 objective lens. **f**, Transgenic animal co-expressing a full-length *p_{dct-1}*:DCT-1::GFP reporter fusion together with *p_{myo-3}*:INVOM::RFP outer mitochondrial membrane reporter. DCT-1::GFP co-localizes with INVOM::RFP, demonstrating that DCT-1 is an outer mitochondrial membrane protein. Scale bars, 10 μm and 5 μm. Images were acquired using a × 100 objective lens. **g**, Isolated mitochondria from transgenic animals expressing a full-length *p_{dct-1}*:DCT-1::GFP reporter fusion were untreated (lane 1, M), treated with proteinase K (lane 2, M^{PK}), or treated with sodium carbonate (lanes 3 and 4) to separate soluble from membrane material. DCT-1::GFP is accessible and degraded by proteinase K while HSP-60, a mitochondrial matrix protein, is protected. Soluble (lane 3, S) and membrane (lane 4, P) fractions after carbonate extraction are separated by ultracentrifugation. DCT-1::GFP is detected in the membrane fraction (P) while soluble HSP-60 protein is found in the soluble material, indicating that DCT-1 is an integral membrane protein.



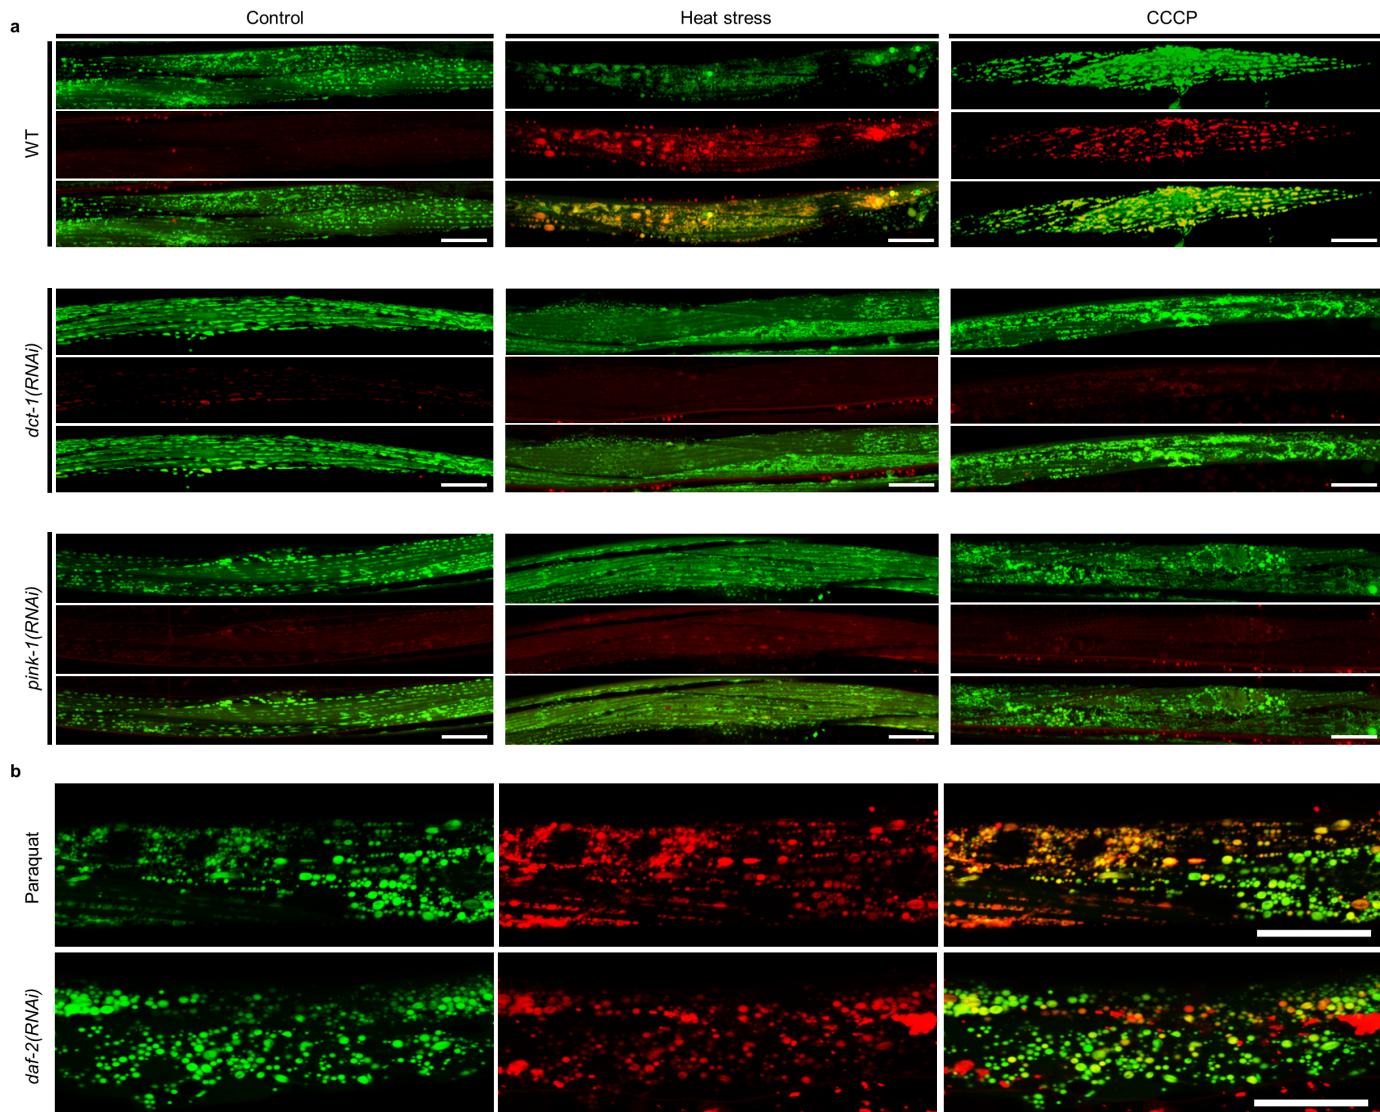
Extended Data Figure 3 | Mitophagy inhibition does not alter intestinal or muscle cell mitochondrial network morphology of FZO-1- and DRP-1-depleted animals. **a**, Mitophagy inhibition by RNAi against *dct-1*, *pink-1* and *pdr-1* does not affect mitochondrial network morphology in FZO-1- and DRP-1-depleted animals that express a mitochondria-targeted GFP in the intestine. The mitochondrial network in wild-type animals is well-organized and highly connected. By contrast, animals subjected to RNAi against *dct-1*, *pink-1*, *pdr-1* and *fzo-1* display fragmented and disorganized mitochondrial network morphology. Simultaneous knocking down of *fzo-1/dct-1*, *pink-1/fzo-1* and *fzo-1/pdr-1* does not alter the already fragmented mitochondrial network of FZO-1-deficient worms. DRP-1-depleted animals display disorganized aggregated and globular mitochondria compared with wild-type worms. Simultaneous knocking down of *drp-1/dct-1*, *pink-1/drp-1* and *pdr-1/drp-1* does not change the impaired mitochondrial network of DRP-1-depleted animals. Scale bar, 20 μ m. Images were acquired using a $\times 40$ objective lens. **b**, Transgenic animals expressing a mitochondria-targeted GFP in muscle cells subjected to RNAi against *dct-1*, *pink-1*, *pdr-1* and *fzo-1*. The mitochondrial network in wild-type animals is well organized and runs parallel with the myofilament lattice. By contrast, DCT-1, PINK-1, PDR-1 and FZO-1 depleted animals display fragmented and disorganized mitochondrial network morphology. Simultaneous knocking down of *fzo-1/dct-1*, *pink-1/fzo-1* and *fzo-1/pdr-1* does not alter the already fragmented mitochondrial network of FZO-1-deficient worms ($n = 120$; *** $P < 0.001$ and NS, $P > 0.05$; one-way ANOVA). **c**, DRP-1-deficient animals present disorganized and aggregated mitochondrial morphology. Simultaneous knocking down of *drp-1/dct-1*, *pink-1/drp-1* and *pdr-1/drp-1* does not change the impaired mitochondrial network of DRP-1-depleted animals ($n = 100$; *** $P < 0.001$ and NS, $P > 0.05$; one-way ANOVA). Error bars, s.e.m.

20 μ m. Images were acquired using a $\times 40$ objective lens. **b**, Transgenic animals expressing a mitochondria-targeted GFP in muscle cells subjected to RNAi against *dct-1*, *pink-1*, *pdr-1* and *fzo-1*. The mitochondrial network in wild-type animals is well organized and runs parallel with the myofilament lattice. By contrast, DCT-1, PINK-1, PDR-1 and FZO-1 depleted animals display fragmented and disorganized mitochondrial network morphology. Simultaneous knocking down of *fzo-1/dct-1*, *pink-1/fzo-1* and *fzo-1/pdr-1* does not alter the already fragmented mitochondrial network of FZO-1-deficient worms ($n = 120$; *** $P < 0.001$ and NS, $P > 0.05$; one-way ANOVA). **c**, DRP-1-deficient animals present disorganized and aggregated mitochondrial morphology. Simultaneous knocking down of *drp-1/dct-1*, *pink-1/drp-1* and *pdr-1/drp-1* does not change the impaired mitochondrial network of DRP-1-depleted animals ($n = 100$; *** $P < 0.001$ and NS, $P > 0.05$; one-way ANOVA). Error bars, s.e.m.



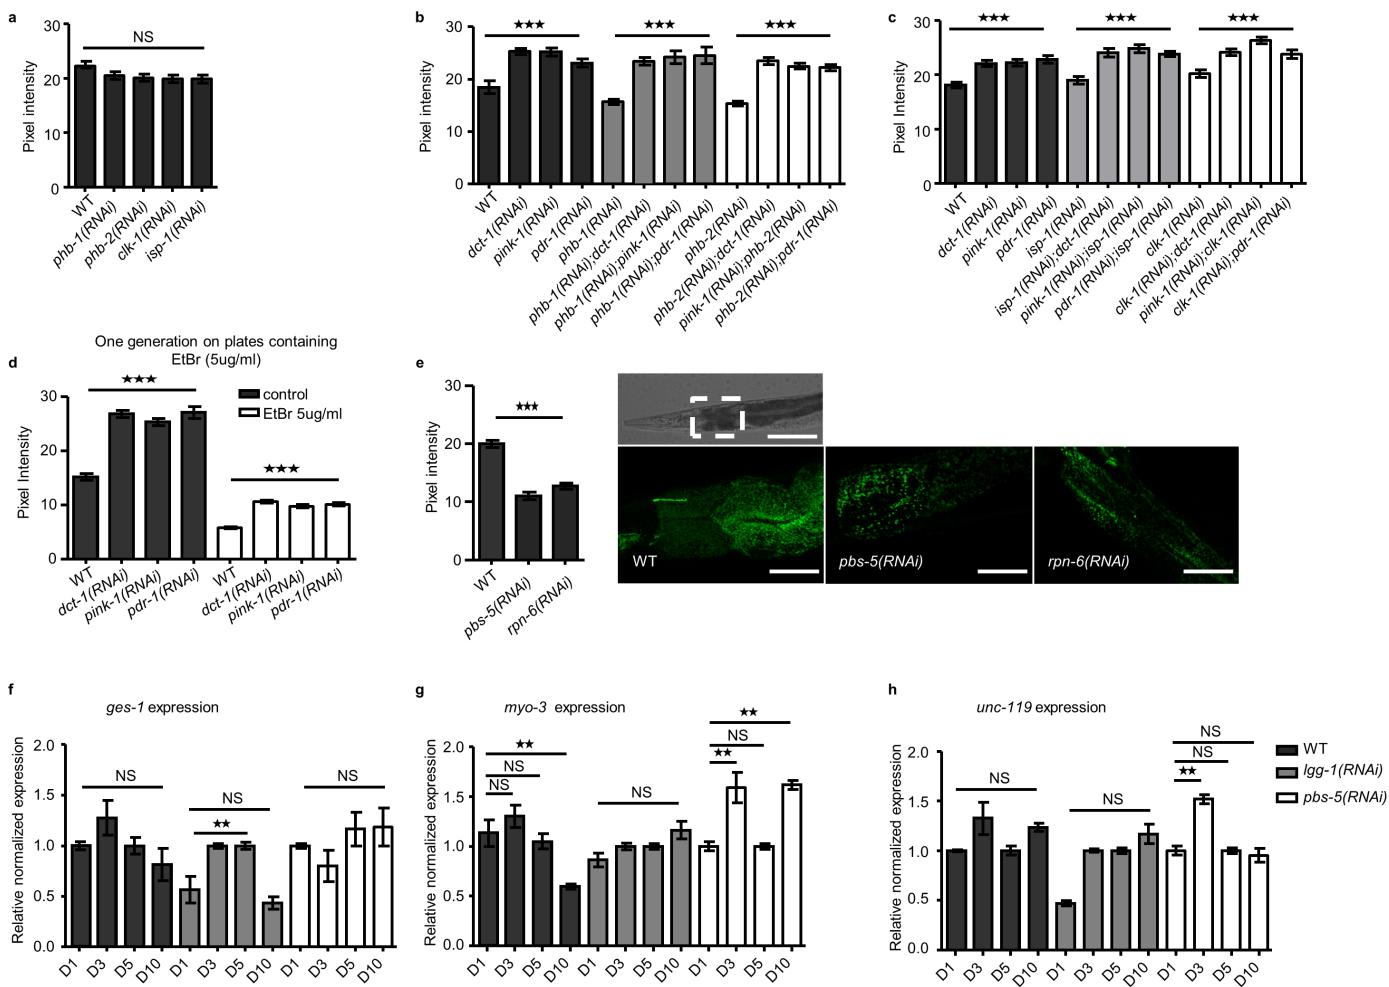
Extended Data Figure 4 | Mitophagy is induced under several stress conditions. **a**, DCT-1 specifically localizes on the outer mitochondrial membrane (shown for muscle cells), where it co-localizes with the autophagosomal protein LGG-1/Atg8. Confocal images of body wall muscle cells of wild-type animals expressing fluorescently tagged DCT-1::GFP along with the autophagosomal marker LGG-1/Atg8 (tagged with DsRed). **b**, DCT-1 co-localizes with the E3 ubiquitin ligase PDR-1/Parkin. Confocal images of body wall muscle cells in wild-type animals expressing fluorescently tagged DCT-1::GFP and PDR-1::DsRed. Scale bar, 20 μ m. Images were acquired using a $\times 40$ objective lens. **c**, *C. elegans* transgenic embryos expressing a full-length $p_{\text{lgg-1}}\text{GFP}:\text{LGG-1}$ fusion protein indicative of autophagic activity subjected to single *dct-1*, *pink-1* RNAi or subjected simultaneous *daf-2/dct-1* and *pink-1/daf-2* RNAi ($n = 70$; NS, $P > 0.05$, *** $P < 0.001$; one-way ANOVA). **d**, *C. elegans* transgenic animals expressing a full-length $p_{\text{lgg-1}}\text{DsRed}:\text{LGG-1}$ fusion protein indicative of autophagic activity subjected to single *daf-2*, *dct-1*, *pink-1*, *pdr-1* RNAi or subjected simultaneous *daf-2/dct-1*, *pink-1/daf-2*, *daf-2/pdr-1* RNAi. Autophagosome number monitored in the pharynx of the animals ($n = 100$; NS, $P > 0.05$, *** $P < 0.001$; one-way ANOVA). **e**, The *daf-2(e1370)*

mutant subjected to RNAi against *dct-1*, *pink-1*, *pdr-1* and *lgg-1* or subjected simultaneously against *lgg-1/dct-1*, *pink-1/lgg-1*, *lgg-1/pdr-1* RNAi. Lifespan values are given in Supplementary Table 1; assays were performed at 20 °C. **f**, Transgenic animals expressing the mtRosella biosensor in body wall muscle cells, were subjected to *daf-2(RNAi)*, treated with paraquat or CCCP and exposed to heat stress (37 °C). Mitophagy stimulation is signified by the reduction of the ratio between pH-sensitive GFP to pH-insensitive DsRed. DCT-1, PINK-1 and PDR-1 are required for mitophagy. Mitophagy is not activated under stress conditions when transgenic animals are subjected to RNAi against *dct-1*, *pink-1* and *pdr-1* ($n = 120$; NS, $P > 0.05$, *** $P < 0.001$; one-way ANOVA). Scale bars, 20 μ m. Images were acquired using a $\times 10$ objective lens. **g**, Mitophagy is not further reduced under oxidative stress conditions in transgenic animals subjected to RNAi against *dct-1*, *pdr-1* or subjected to simultaneous *pdr-1/dct-1* RNAi (NS, $P > 0.05$, *** $P < 0.001$; one-way ANOVA). **h**, Mitophagy inhibition under stress conditions in transgenic animals subjected to RNAi against *lgg-1* (NS, $P > 0.05$, *** $P < 0.001$; one-way ANOVA). Error bars, s.e.m.



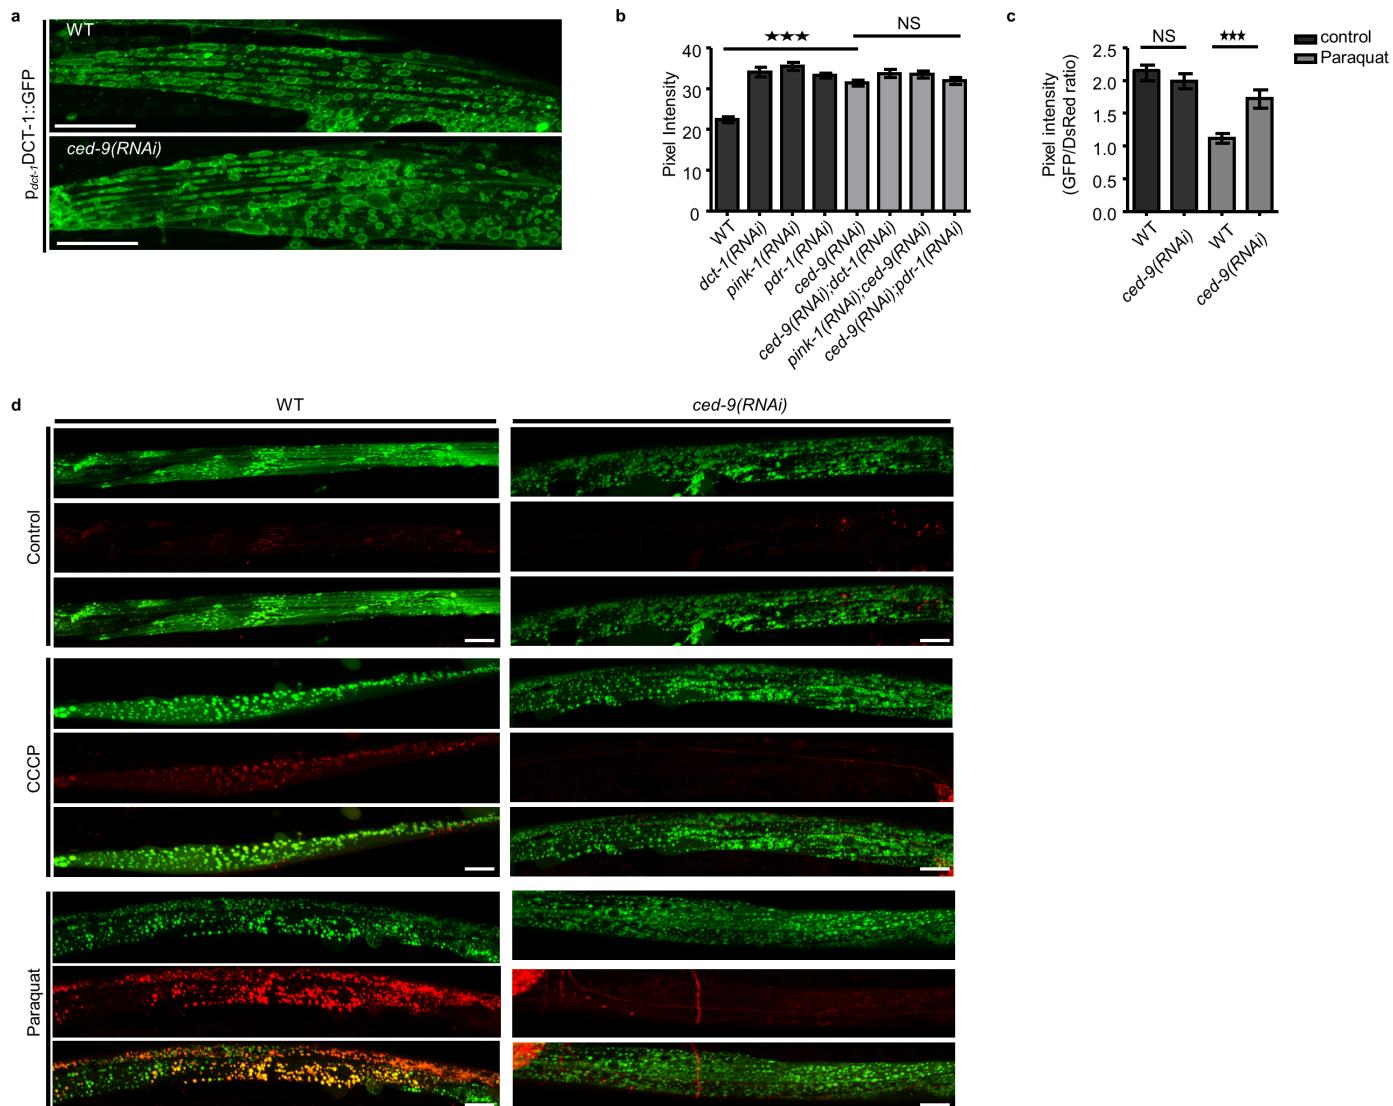
Extended Data Figure 5 | Monitoring mitophagy *in vivo*. Transgenic animals co-expressing a mitochondria-tagged GFP (mtGFP) in body wall muscle cells and the autophagosomal protein LGG-1 fused with DsRed, were exposed to heat stress (37 °C), treated with CCCP, paraquat or subjected to RNAi against *daf-2*. Mitophagy induction is signified by co-localization of GFP and DsRed signals (for each group of images mitochondria are shown in green

on top, autophagosomes in red below, with a merged image at the bottom). **a**, DCT-1 and PINK-1 are required for mitophagy. Mitophagy inhibition under stress conditions when transgenic animals are subjected to RNAi against *dct-1* and *pink-1*. Scale bars, 20 µm. Images were acquired using a ×40 objective lens. **b**, High-magnification images showing induction of mitophagy. Scale bars, 10 µm. Images were acquired using a ×40 objective lens.



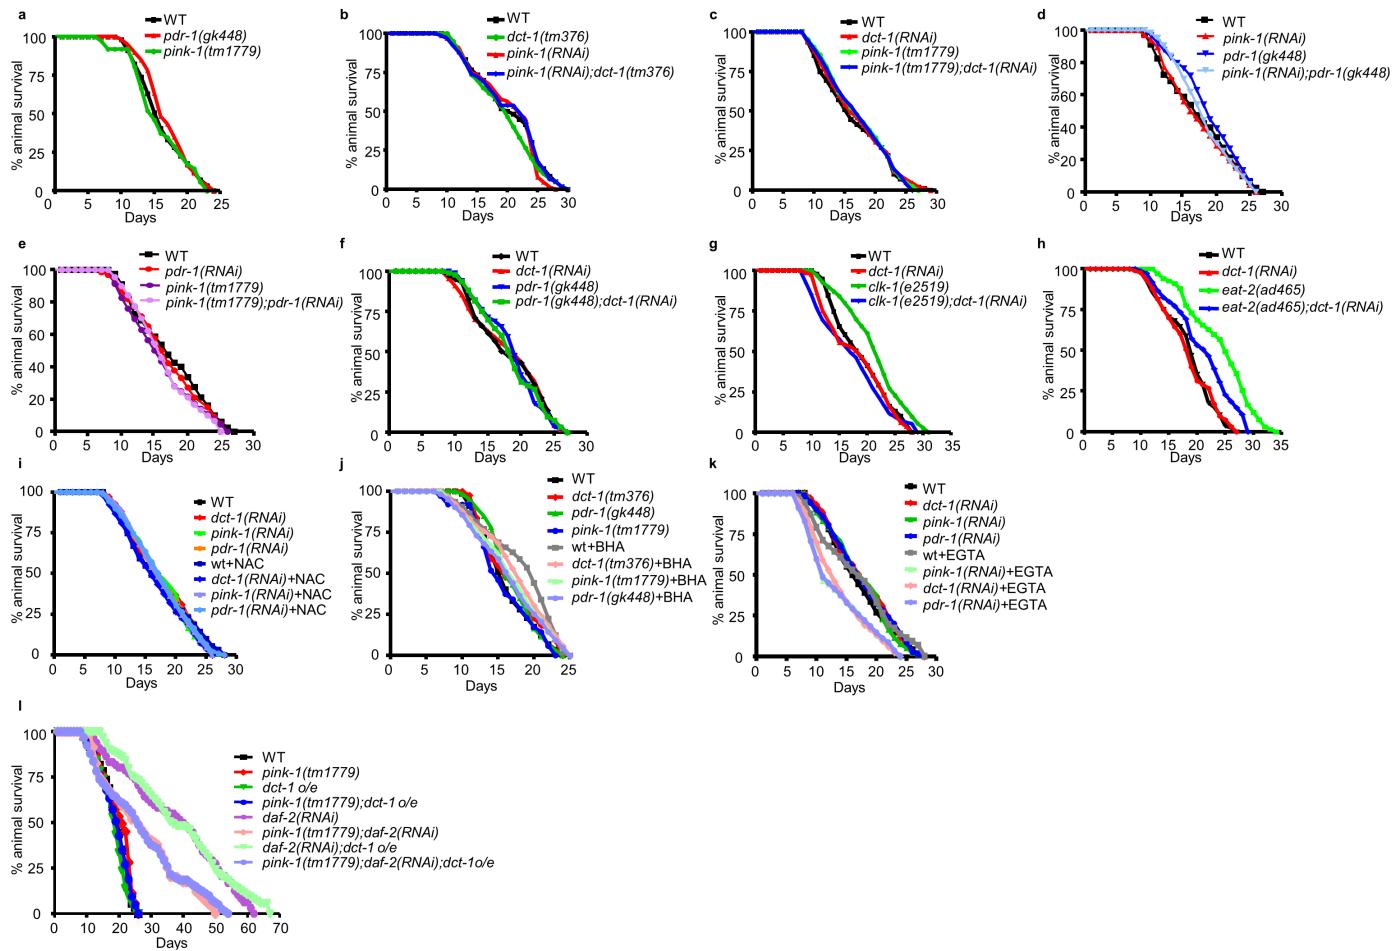
Extended Data Figure 6 | Mitochondrial content is not affected upon mitochondrial dysfunction. **a**, Knockdown of either PHB-1, PHB-2, CLK-1 or ISP-1 does not affect the mitochondrial mass in transgenic worms that express a mitochondria-targeted GFP in the intestine ($n = 100$; NS, $P > 0.05$; one-way ANOVA). **b, c**, Simultaneous RNAi against *phb-1/dct-1*, *phb-1/pink-1*, *phb-1/pdr-1*, *phb-2/dct-1*, *pink-1/phb-2*, *phb-2/pdr-1*, *isp-1/dct-1*, *pink-1/isp-1*, *pdr-1/isp-1*, *clk-1/dct-1*, *pink-1/clk-1*, *clk-1/pdr-1* increases mitochondrial mass ($n = 100$; *** $P < 0.001$; one-way ANOVA). **d**, Knockdown of either DCT-1, PINK-1 or PDR-1 increases mitochondrial mass in transgenic worms that express a mitochondria-targeted GFP in the intestine. Transgenic worms expressing mitochondria-targeted GFP in the intestine were raised over one

generation in the presence of EtBr to block mitochondrial biogenesis. Knockdown of either DCT-1, PINK-1 or PDR-1 increases mitochondrial content of animals exposed to EtBr ($n = 100$; *** $P < 0.001$; one-way ANOVA). **e**, Knockdown of either PBS-5 or RPN-6 promotes fragmentation of mitochondrial network and decreases mitochondrial content in worms that express a mitochondria-targeted GFP in the intestine ($n = 100$; *** $P < 0.001$; one-way ANOVA). Scale bar, 20 μ m. Images were acquired using a $\times 40$ objective lens. **f-h**, The expression levels of *ges-1*, *myo-3* and *unc-119* gene are relatively stable throughout ageing. Inhibition of autophagy or proteasome system does not alter the expression levels of *ges-1*, *myo-3* and *unc-119* genes during ageing (NS, $P > 0.05$, ** $P < 0.01$; one-way ANOVA). Error bars, s.e.m.



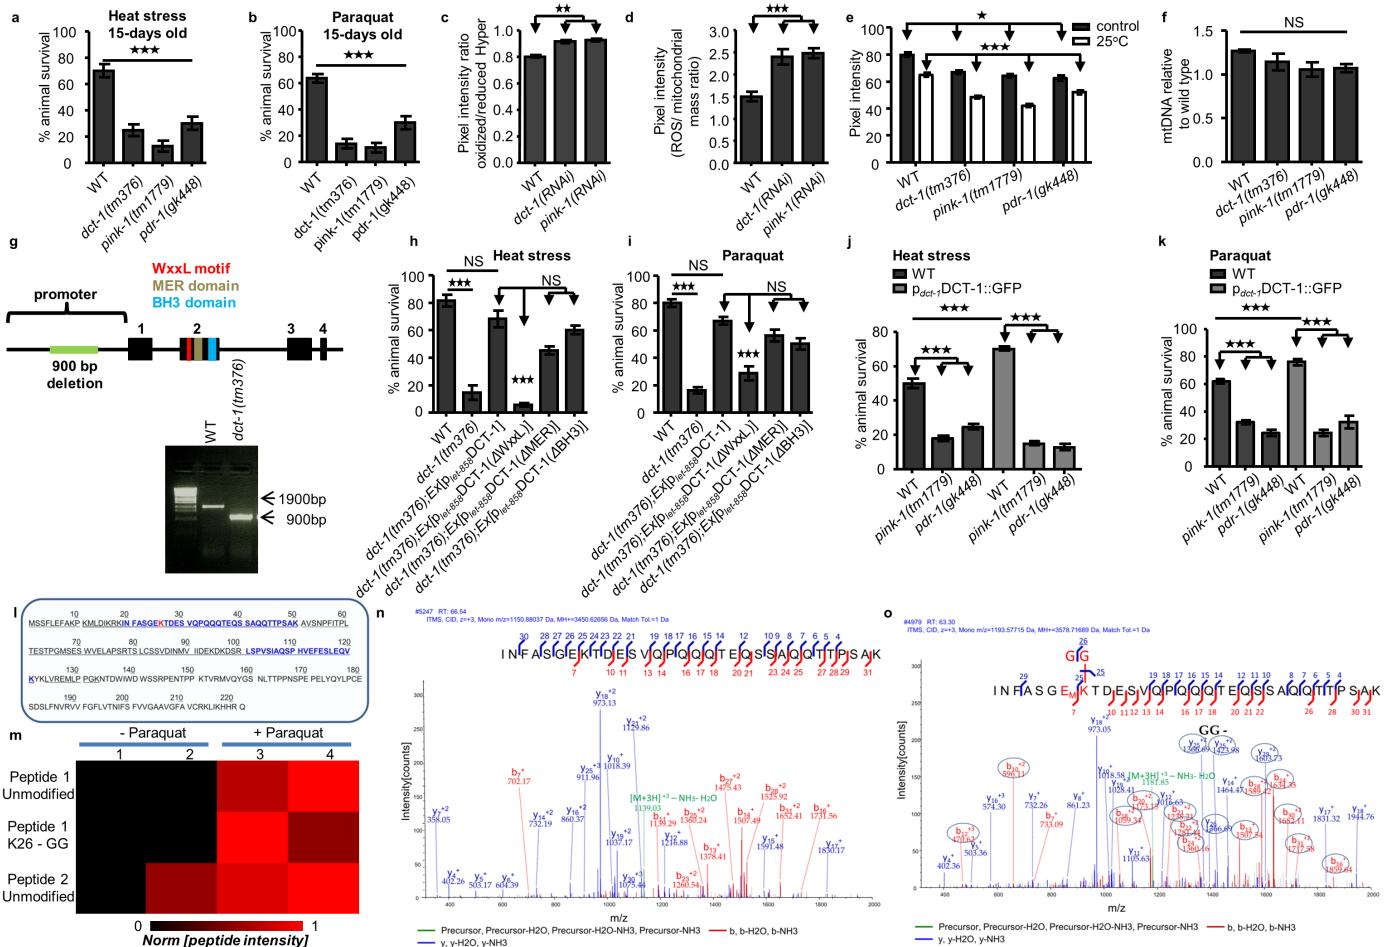
Extended Data Figure 7 | Gene ced-9 acts in the same genetic pathway with *dct-1*, *pink-1* and *pdr-1* to mediate mitophagy. **a**, CED-9 depletion does not affect the localization of DCT-1. Scale bars, 20 μ m. Images were acquired using a $\times 63$ objective lens. **b**, Knockdown of either CED-9, DCT-1, PINK-1 or PDR-1 increases mitochondrial mass in transgenic worms expressing a mitochondria-targeted GFP in the intestine. Simultaneous RNAi against *ced-9*/*dct-1*, *pink-1*/*ced-9* and *ced-9*/*pdr-1* does not increase further mitochondrial mass ($n = 100$; NS, $P > 0.05$, *** $P < 0.001$; one-way ANOVA). **c**, Transgenic

animals expressing the mtRosella biosensor in body wall muscle cells, were treated with paraquat. Mitophagy induction is signified by the reduction of the ratio between pH-sensitive GFP to pH-insensitive DsRed. CED-9 is required for mitophagy. Mitophagy inhibition under stress conditions when transgenic animals are subjected to RNAi against *ced-9* ($n = 120$; NS, $P > 0.05$, *** $P < 0.001$; one-way ANOVA). Error bars, s.e.m. **d**, Mitophagy inhibition under stress, in transgenic animals subjected to RNAi against *ced-9*. Scale bars, 20 μ m. Images were acquired using a $\times 40$ objective lens.



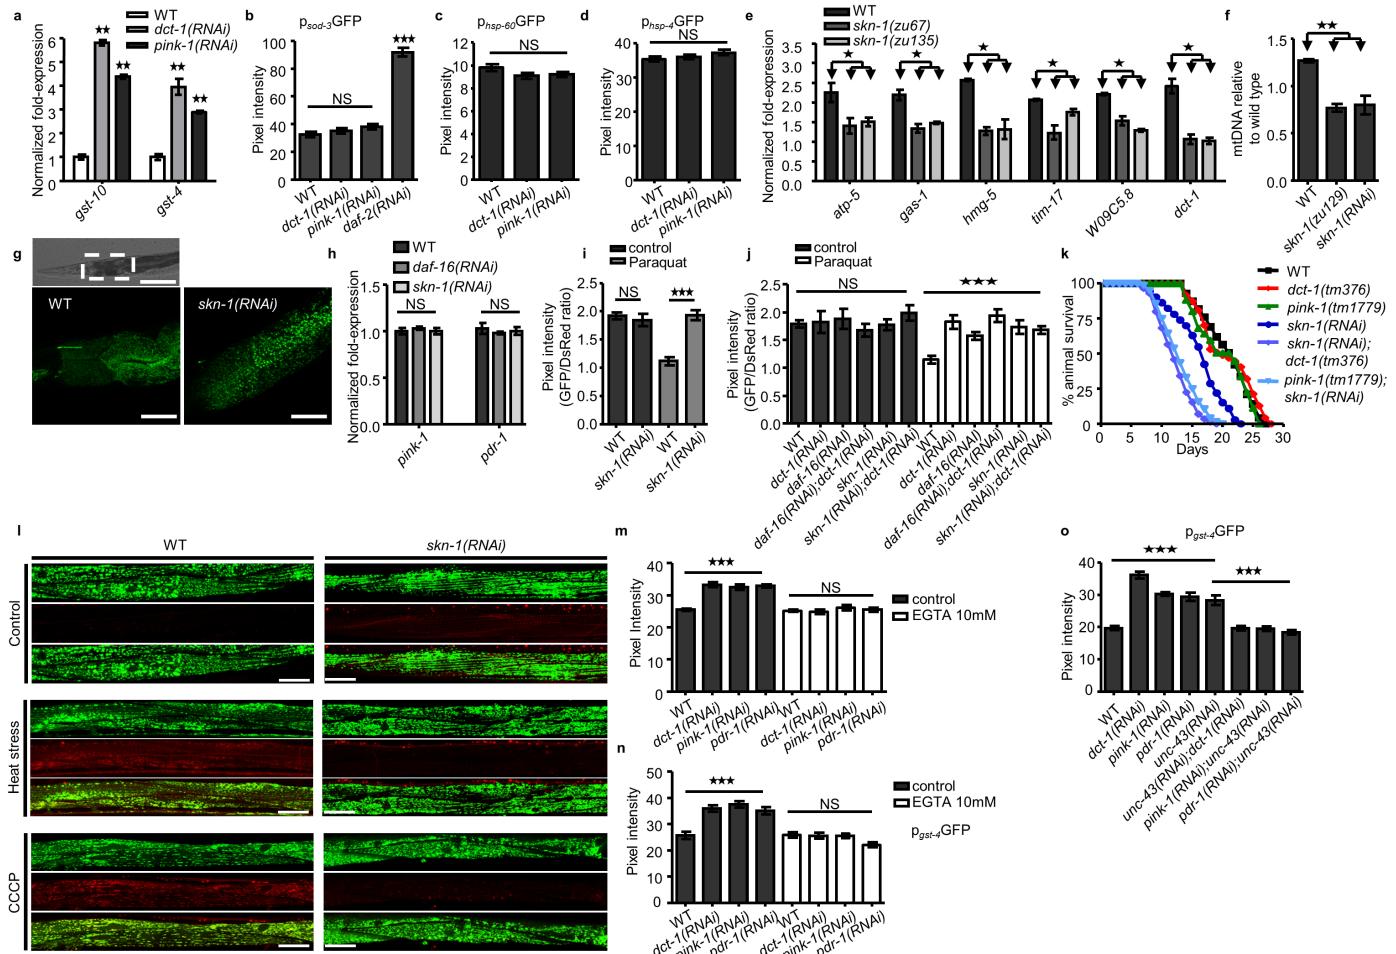
Extended Data Figure 8 | Inhibition of mitophagy affects longevity in a genetic-background-specific manner. **a**, Depletion of PDR-1, or PINK-1. **b, c**, Depletion of both DCT-1 and PINK-1. **d, e**, Depletion of both PDR-1 and PINK-1. **f**, Depletion of both DCT-1 and PDR-1. Knockdown of DCT-1 shortens the lifespan of long-lived *clk-1(e2519)* (**g**), and *eat-2(ad465)* (**h**) mutants. **i, j**, Mitochormesis is not engaged in mitophagy-deficient animals to influence lifespan. Depletion of either DCT-1, PINK-1 or PDR-1, with, or without NAC or BHA treatment. **k**, Depletion of DCT-1, PINK-1 or PDR-1, with, or without EGTA treatment. **l**, DCT-1 overexpression does not influence the lifespan of otherwise wild-type animals or *pink-1(tm1779)* mutants under normal conditions. DCT-1 overexpression extends only the maximum lifespan of DAF-2-depleted worms. However, PINK-1 deficiency abolishes the extended lifespan of both *daf-2(RNAi)* and *daf-2(RNAi);Ex[p_{dct-1}DCT-1:GFP]* animals. Survival curves depict the percentage of animals remaining alive, plotted against animal age. Lifespan values are given in Supplementary Table 1; assays were performed at 20 °C.

with, or without EGTA treatment. **l**, DCT-1 overexpression does not influence the lifespan of otherwise wild-type animals or *pink-1(tm1779)* mutants under normal conditions. DCT-1 overexpression extends only the maximum lifespan of DAF-2-depleted worms. However, PINK-1 deficiency abolishes the extended lifespan of both *daf-2(RNAi)* and *daf-2(RNAi);Ex[p_{dct-1}DCT-1:GFP]* animals. Survival curves depict the percentage of animals remaining alive, plotted against animal age. Lifespan values are given in Supplementary Table 1; assays were performed at 20 °C.



Extended Data Figure 9 | Mitophagy deficiency impairs mitochondrial homeostasis and stress resistance. **a, b**, Fifteen-day-old *dct-1(tm376)*, *pink-1(tm1779)*, *pdr-1(gk448)* and wild-type animals incubated (**a**) at 37 °C for 7 h ($n = 150$; *** $P < 0.001$; one-way ANOVA) and (**b**) treated with 8 mM paraquat ($n = 150$; *** $P < 0.001$; one-way ANOVA). **c**, Transgenic animals expressing the H₂O₂-biosensor HyPer subjected to either *dct-1* or *pink-1* RNAi. The graph depicts pixel intensity ratios of oxidized to reduced HyPer ($n = 70$; ** $P < 0.01$; one-way ANOVA). **d**, Mitophagy-depleted animals display increased ratios of mitochondrial ROS formation (MitoTracker Red CM-H₂X ROS) to total mitochondrial content (intestinal mtGFP) ($n = 150$; *** $P < 0.001$; one-way ANOVA). **e**, TMRE staining of mitophagy-depleted adults at 20 °C or 25 °C ($n = 150$; * $P < 0.05$, *** $P < 0.001$; one-way ANOVA). **f**, The percentage of mitochondrial DNA relative to wild type is shown. Quantitative PCR was performed in wild-type, *dct-1(tm376)*, *pink-1(tm1779)* and *pdr-1(gk448)* mutant animals. The percentage of mtDNA is not affected in mitophagy-deficient animals (NS, $P > 0.05$; one-way ANOVA). **g**, The genetic locus of *dct-1* gene is depicted; *dct-1(tm376)* is homozygous deletion. The *dct-1(tm376)* mutant carries a deletion of 900 base pairs at the promoter region of *dct-1* gene. **h**, **i**, Transgenic animals were generated carrying extra copies of wild-type *dct-1* gene and *dct-1* gene carrying deletions for the corresponding sequences of WXXL motif, MER domain and BH3 domain, and were exposed to heat stress (37 °C) or treated with paraquat. WXXL motif of DCT-1 is required for mitophagy process since WXXL-deleted DCT-1 fails to rescue resistance of *dct-1(tm376)* animals under stress conditions ($n = 150$; NS, $P > 0.05$, *** $P < 0.001$; one-way ANOVA). **j**, **k**, Wild-type, *pink-1(tm1779)* and *pdr-1(gk448)* transgenic animals expressing a full-length *p_{dct-1}DCT-1::GFP* reporter fusion were exposed to heat stress (37 °C) or treated with paraquat. DCT-1 overexpression is sufficient to promote stress resistance in wild-type worms. However, the stress resistance of DCT-1 overexpressing animals is abolished in *pink-1* and *pdr-1* mutant background ($n = 150$; *** $P < 0.001$; one-way ANOVA). Error bars, s.e.m. **l-o**, DCT-1 is ubiquitinated in a PINK-1-dependent manner upon oxidative stress. DCT-1::GFP was

immunoprecipitated from total lysates of wild-type and *pink-1(tm1779)* transgenic animals expressing a full-length *p_{dct-1}DCT-1::GFP* reporter fusion. The immunoprecipitated material was trypsinized and subjected to nLC-MS/MS analysis. **l**, DCT-1 was identified by mass spectrometry with high confidence and a sequence coverage of ~60% (underlined on DCT-1 sequence). The two most prominent tryptic peptides of DCT-1 (highlighted in blue) were used for the relative quantification of the ubiquitination within the different samples. Ubiquitin-modified lysine is highlighted in red. **m**, Numbers represent the following samples: untreated wild-type worms (sample 1), untreated wild-type transgenic worms expressing a full-length *p_{dct-1}DCT-1::GFP* reporter fusion (sample 2), paraquat-treated wild-type transgenic worms expressing a full-length *p_{dct-1}DCT-1::GFP* (sample 3) and paraquat-treated *pink-1(tm1779)* mutants expressing a full-length *p_{dct-1}DCT-1::GFP* (sample 4). The most prominently identified tryptic peptides 19–50 (peptide 1) and 101–121 (peptide 2) were used to determine the relative abundance of modified and unmodified DCT-1 in the four samples. The maximum intensities of the peptide ions of each peptide in the different samples were normalized and used for the comparative analysis and K26-GG abundance. The unmodified peptides showed similar pattern and abundance in the different samples (not identified in control sample 1, which did not express DCT-1::GFP, as expected), while the tryptic peptide with the K26-GG modification (19–50) was significantly enriched in sample 3 where mitophagy was induced in the presence of PINK-1. **n**, **o**, K26-GG modification on the tryptic peptide 19–50 of DCT-1 was manually validated on the MS/MS spectra of the unmodified (**n**) and modified (**o**) peptide. Both peptides were identified by high amino-acid coverage, as is shown on their corresponding assigned spectra. The intact peptide masses after the loss of H₂O and NH₃ molecules are shown in green, and the characteristic *b* and *y* fragment ions used for sequencing the amino-acid residues of the peptides are shown in red and blue, respectively. The characteristic fragments indicating the K26-GG modification are marked by blue circles and underlined fragment numbers (**o**).



Extended Data Figure 10 | SKN-1 regulates the transcription of mitochondrial genes and modulates mitochondrial homeostasis and integrity. **a**, mRNA level analysis of the SKN-1 target genes *gst-4* and *gst-10* in wild-type animals and mitophagy-deficient animals (**P < 0.01; one-way ANOVA). **b**, Expression of the DAF-16-regulated, p_{sod-3}GFP transgene in animals subjected to *dct-1* or *pink-1* RNAi (n = 70; NS, P > 0.05, ***P < 0.001; one-way ANOVA). **c**, Expression of the mitochondrial stress reporter p_{hsp-60}GFP, in animals subjected to RNAi against *dct-1* or *pink-1* (n = 70; NS, P > 0.05; one-way ANOVA). **d**, Expression of the ER stress reporter p_{hsp-4}GFP, in animals subjected to RNAi against *dct-1* or *pink-1* (n = 70; NS, P > 0.05; one-way ANOVA). **e**, mRNA levels of *atp-5*, *gas-1*, *hmg-5*, *tim-17*, W09C5.8 and *dct-1* in wild-type and SKN-1-depleted animals (*skn-1(zu67)* and *skn-1(zu135)* animals) (*P < 0.01; one-way ANOVA). **f**, The percentage of mitochondrial DNA relative to wild type is shown. The percentage of mtDNA is reduced in SKN-1-depleted worms (**P < 0.01; one-way ANOVA). **g**, Knockdown of SKN-1 alters mitochondrial network morphology in the intestine of young adult animals. Scale bar, 20 μm. Images were acquired using a ×40 objective lens. **h**, mRNA levels of *pink-1* and *pdr-1* in wild-type animals, DAF-16- and SKN-1-depleted animals (NS, P > 0.05; one-way ANOVA). **i**, Mitophagy is not stimulated under oxidative stress, in transgenic animals expressing mtRosella subjected to RNAi against *skn-1* (n = 120; NS, P > 0.05, ***P < 0.001; one-way ANOVA). **j**, Monitoring mitophagy under oxidative stress, in transgenic

animals subjected to RNAi against *skn-1*, *skn-1* and *dct-1*, *daf-16*, *daf-16* and *dct-1* (n = 120; NS, P > 0.05, ***P < 0.001; one-way ANOVA). **k**, SKN-1 compensates for inhibition of mitophagy in otherwise wild-type animals. Simultaneous depletion of DCT-1 and SKN-1, or PINK-1 and SKN-1. Survival curves depict the percentage of animals remaining alive, plotted against animal age. Lifespan values are given in Supplementary Table 1; assays were performed at 20 °C. **l**, Transgenic animals co-expressing a mitochondria-tagged GFP (mtGFP) in body wall muscle cells and the autophagosomal protein LGG-1 fused to DsRed were exposed to heat stress (37 °C) or treated with CCCP. Induction of mitophagy, as indicated by co-localization of GFP and DsRed (for each group of images, mitochondria are shown in green on top, autophagosomes in red below, with a merged image at the bottom). Mitophagy inhibition under stress, in transgenic animals subjected to RNAi against *skn-1*. Scale bars, 20 μm. Images were acquired using a ×40 objective lens. **m**, Cytoplasmic calcium elevation upon mitophagy depletion is abolished upon treatment with EGTA (n = 100; NS, P > 0.05, ***P < 0.001; one-way ANOVA). **n**, Transgenic animals expressing p_{gst-4}GFP reporter subjected to RNAi against *dct-1*, *pink-1* and *pdr-1*, with or without EGTA treatment (n = 120; NS, P > 0.05, ***P < 0.001; one-way ANOVA). **o**, UNC-43 regulates the transcriptional activity of SKN-1 upon mitophagy depletion (n = 90; ***P < 0.001; one-way ANOVA). Error bars, s.e.m.

Coordinated regulation of bidirectional COPI transport at the Golgi by CDC42

Seung-Yeol Park^{1,2}, Jia-Shu Yang^{1,2}, Angela B. Schmider^{2,3}, Roy J. Soberman^{2,3,4} & Victor W. Hsu^{1,2}

The Golgi complex has a central role in the intracellular sorting of secretory proteins^{1,2}. Anterograde transport through the Golgi has been explained by the movement of Golgi cisternae, known as cisternal maturation^{3–5}. Because this explanation is now appreciated to be incomplete⁶, interest has developed in understanding tubules that connect the Golgi cisternae^{7–9}. Here we show that the coat protein I (COPI) complex sorts anterograde cargoes into these tubules in human cells. Moreover, the small GTPase CDC42 regulates bidirectional Golgi transport by targeting the dual functions of COPI in cargo sorting and carrier formation. CDC42 also directly imparts membrane curvature to promote COPI tubule formation. Our findings further reveal that COPI tubular transport complements cisternal maturation in explaining how anterograde Golgi transport is achieved, and that bidirectional COPI transport is modulated by environmental cues through CDC42.

In addition to its known role in generating vesicles¹⁰, we recently discovered that COPI also generates tubules that connect the Golgi stacks⁹. However, whether these tubules act passively or actively in cargo transport has been unclear. Active transport involves coat proteins binding to cargoes. Thus, we initially examined whether COPI binds to a temperature-sensitive form of the vesicular stomatitis virus G protein (hereafter referred as VSVG), which has been widely used to track anterograde Golgi transport^{6–9}. We found that coatomer, the core component of the COPI complex¹⁰, binds directly to the cytoplasmic tail of VSVG (Fig. 1a and Extended Data Fig. 1a). Further defining this binding, we found that coatomer binds to the membrane-proximal region of the VSVG tail (Fig. 1b and Extended Data Fig. 1b), and then identified residues within this region crucial for binding by coatomer (Fig. 1c and Extended Data Fig. 1a).

Coatomer has been shown previously to promote the retrograde transport of cargoes, such as WBP1 (ref. 11) and the KDEL receptor¹², which involves binding to their carboxy terminus that contains di-lysine residues. By contrast, the basic residues in VSVG recognized by coatomer are situated away from this end (see Extended Data Fig. 1a), suggesting a new mode of cargo recognition by coatomer. As confirmation, we performed competition studies, and verified that a peptide derived from WBP1 (containing the retrograde di-lysine motif) cannot compete with the VSVG tail for binding to coatomer (Fig. 1d). We also used fluorescence lifetime imaging microscopy (FLIM) to confirm that coatomer interacts with VSVG at the Golgi, and the interaction requires crucial basic residues in VSVG (Fig. 1e and Extended Data Fig. 1c). Subsequently, we found that mutation of these basic residues delayed the transport of VSVG from the endoplasmic reticulum (ER) to the *trans*-Golgi (Fig. 1f and Extended Data Fig. 1d), but not from the ER to the *cis*-Golgi (Extended Data Fig. 1e). Thus, coatomer promotes the intra-Golgi transport of VSVG through cargo binding.

We next examined another anterograde cargo, the endogenous low-density lipoprotein receptor (LDLR). Coatomer also binds directly to its cytoplasmic tail (Fig. 2a), which is also not affected by the presence

of the WBP1 peptide (Fig. 2b). As the LDLR tail also contains basic residues away from the C terminus (Extended Data Fig. 2a), we next targeted them for mutagenesis and found that binding to the LDLR tail by coatomer became reduced (Fig. 2a). To confirm the functional role of this binding, we swapped the tails of LDLR and VSVG, and then examined the transport of the resulting chimaera (VSVG-LDLR). When crucial basic residues in the LDLR tail were mutated, we observed delayed transport of VSVG-LDLR from the ER to the *trans*-Golgi (Fig. 2c and Extended Data Fig. 2b), but not from the ER to the *cis*-Golgi (Extended Data Fig. 2c).

As compared to VSVG, VSVG-LDLR is transported slower from the ER to the *cis*-Golgi (Extended Data Figs 1e and 2c), and also through the Golgi (Fig. 2d and Extended Data Fig. 2d). Pursuing the latter observation, we found that coatomer binds more efficiently to the VSVG tail than the LDLR tail, both *in vitro* (Fig. 2e) and in cells (Fig. 2f and Extended Data Fig. 3a). Moreover, mutating critical basic residues in either cargo tail results in transport rates becoming similar (Fig. 2d and Extended Data Fig. 2d). Furthermore, mutating a residue in VSVG that partially reduces its binding by coatomer (His494Ala in Fig. 1c) moderately reduces VSVG transport through the Golgi (Extended Data Fig. 3b). Thus, the collective results further confirmed that binding of anterograde cargoes by coatomer promotes their transport through the Golgi.

We then sought to address how the same coat complex can promote opposing directions of Golgi transport while also maintaining specificity in cargo sorting. As a clue, we noted that CDC42 has been found previously to interact with coatomer¹³. Moreover, the CDC42(Phe28Leu) point mutation, which promotes the rapid cycling of its GTPase cycle, enhances VSVG transport¹³. We initially defined more precisely that active CDC42(Phe28Leu) promotes the intra-Golgi segment of VSVG transport (Fig. 3a and Extended Data Fig. 4a), with short interfering RNA (siRNA) against CDC42 (Extended Data Fig. 4b) having the opposite effect (Extended Data Fig. 4c). By contrast, no effect was seen on transport from the ER to the *cis*-Golgi (Extended Data Fig. 4d, e). We then considered that COPI tubular transport at the Golgi requires cytosolic phospholipase A_{2α} (cPLA_{2α}) activity⁹. Targeting this activity using siRNA, we found that CDC42 can no longer promote the anterograde transport of VSVG through the Golgi (Fig. 3b and Extended Data Fig. 5a). Next, examining retrograde COPI vesicular transport, as tracked through a model cargo (VSVG-KDEL^{9,14,15}), we found that active CDC42(Phe28Leu) inhibits this transport (Fig. 3c and Extended Data Fig. 5b), whereas CDC42 siRNA enhances this transport (Extended Data Fig. 5c). Thus, the collective results revealed that CDC42 has opposite effects on the two directions of COPI transport at the Golgi, promoting anterograde tubular transport and inhibiting retrograde vesicular transport.

We then sought to determine how CDC42 exerts these effects. We had previously found that retrograde cargoes are preferentially sorted into COPI vesicles generated through a reconstitution system⁹. However, reconstituted COPI tubules contain both anterograde and

¹Division of Rheumatology, Immunology and Allergy, Brigham and Women's Hospital, Boston, Massachusetts 02115, USA. ²Department of Medicine, Harvard Medical School, Boston, Massachusetts 02115, USA. ³Nephrology Division and Department of Medicine, Massachusetts General Hospital, Charlestown, Massachusetts 02129, USA. ⁴Molecular Imaging Core, Massachusetts General Hospital, Charlestown, Massachusetts 02129, USA.

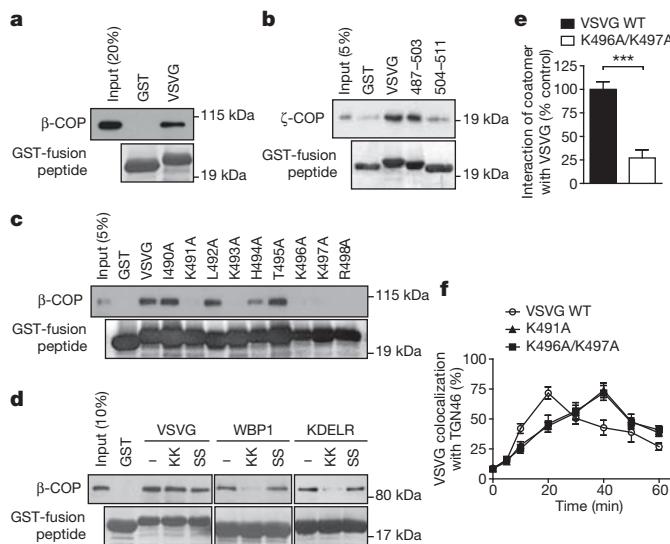


Figure 1 | Coatomer binds VSVG to promote its transport through the Golgi. **a–c**, *In vitro* binding of the VSVG tail (full-length (a), truncations (b), point mutants (c)) by coatomer; $n = 3$. kDa, kilodaltons. **d**, Using the WBP1 peptide that contains the retrograde di-lysine motif (KK) or a di-serine motif (SS) for competition in cargo binding by coatomer; $n = 4$. **e**, FLIM examining VSVG interacting with coatomer; $n = 3$. *** $P < 0.001$ (two-tailed Mann–Whitney test). WT, wild type. **f**, Effect of mutating the VSVG tail on its transport from the ER to the *trans*-Golgi; $n = 4$. Data are mean \pm s.e.m.

retrograde cargoes⁹. Thus, to explore whether CDC42 acts to specify cargo sorting into these tubules, we added CDC42 to the reconstitution system. Remarkably, the level of a retrograde cargo (VSVG-KDELR) in COPI tubules became reduced (Fig. 3d), while that of an anterograde cargo (VSVG) was unaffected (Fig. 3e). In considering how CDC42 achieves these effects, we noted that CDC42 possesses a di-lysine motif, which competes with retrograde cargoes for binding to coatomer¹³. Pursuing this clue, we found that active CDC42 prevents coatomer

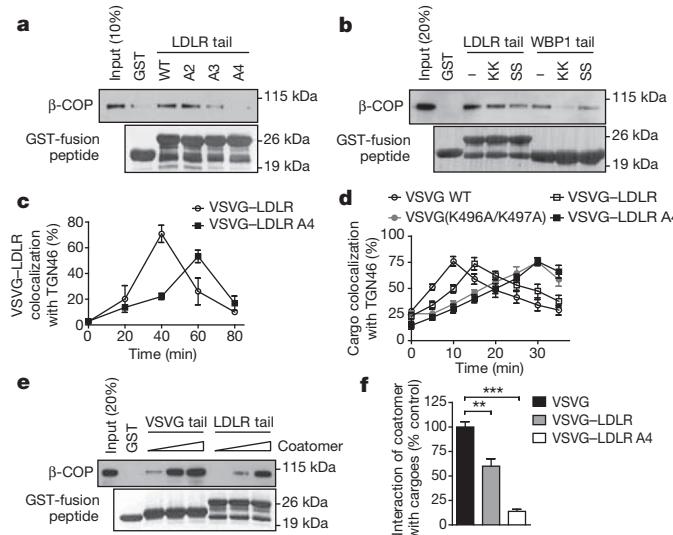


Figure 2 | Coatomer also binds to the LDLR tail to promote the transport of VSVG-LDLR through the Golgi. **a**, *In vitro* binding of the LDLR tail (full-length and point mutants) by coatomer; $n = 4$. **b**, Using the WBP1 peptides for competition in cargo binding by coatomer; $n = 3$. **c**, Effect of mutating the LDLR tail on the transport of VSVG-LDLR from the ER to the *trans*-Golgi; $n = 4$. **d**, Effect of mutations in cargo tails on the intra-Golgi transport of different cargoes; $n = 3$. **e**, Titration of coatomer level for *in vitro* binding to cargo tails; $n = 2$. **f**, FLIM examining different cargoes interacting with coatomer; $n = 3$. *** $P < 0.01$, *** $P < 0.001$ (two-tailed Mann–Whitney test). Data are mean \pm s.e.m.

from binding to the retrograde cargo (KDELR) tail (Fig. 3f), but not coatomer binding to the anterograde cargo (VSVG) tail (Fig. 3g). These results were further confirmed by FLIM (Extended Data Fig. 5d, e). We also examined a range of other small GTPases known to act in Golgi transport, and found that none of them affects coatomer binding to either the VSVG tail (Extended Data Fig. 6a) or the KDELR tail (Extended Data Fig. 6b). Other Rho small GTPases also did not have an effect (Extended Data Fig. 6a, b). Thus, the collective results suggested that CDC42 imparts specificity in cargo sorting into COPI tubules by competing with retrograde, but not anterograde, cargoes for binding to coatomer.

Besides cargo sorting, the other major function of coat proteins is the generation of transport carriers. Previously, we found that the addition of cPLA₂ α to the COPI reconstitution system diverts vesicle formation to tubule formation⁹. Remarkably, when we replaced cPLA₂ α with the active CDC42(Phe28Leu), COPI tubule formation was also promoted (Fig. 4a), and at the expense of vesicle formation (Fig. 4b). In considering how CDC42 exerts these effects, we initially ruled out that CDC42 acts through cPLA₂ α , as the replacement of wild-type cPLA₂ α with the catalytic dead form in the COPI reconstitution system did not prevent active CDC42 from promoting tubule formation (Extended Data Fig. 6c), and also inhibiting vesicle formation (Extended Data Fig. 6d). To decipher a different explanation, we noted that protein factors can directly impart membrane curvature, as reflected by their ability to tubulate liposomes^{16,17}. Notably, when we incubated active CDC42(Phe28Leu) with liposomes, tubulation was observed (Fig. 4c, d).

To determine how CDC42 exerts this effect, we initially considered that CDC42 is prenylated, which suggests how the GTPase can insert asymmetrically into the membrane bilayer to impart curvature. As confirmation, we found that a non-prenylated form of CDC42 cannot induce liposome tubulation (Fig. 4c, d). We then considered that CDC42 undergoes dimerization¹⁸, which suggested an additional possibility that it can polymerize onto membrane to exert curvature through scaffolding. Performing mutagenesis, we identified three residues at the C terminus of CDC42, anchored by di-arginine residues

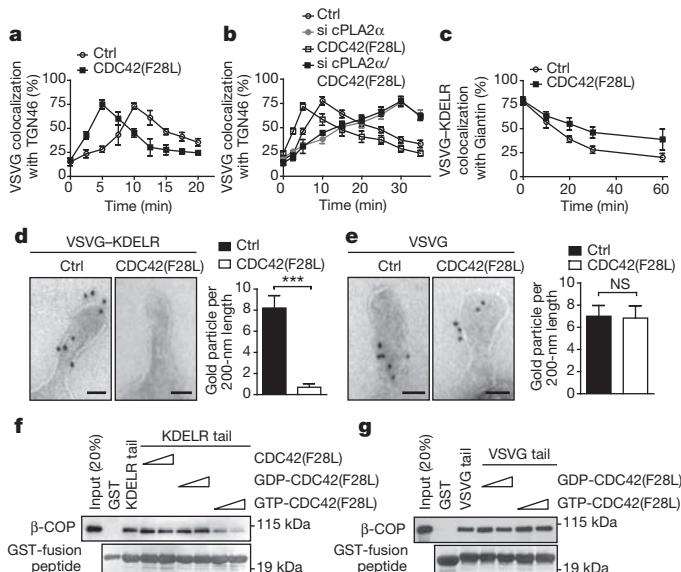


Figure 3 | CDC42 modulates cargo sorting by COPI. **a, b**, Effect of different conditions on the intra-Golgi transport of VSVG; $n = 5$. Ctrl, control; si, siRNA. **c**, Effect of expressing active CDC42 on VSVG-KDELR transport; $n = 3$. **d, e**, Immunogold labelling of reconstituted COPI tubules, with representative (out of 15) electron microscopy images and quantification shown; $n = 4$. *** $P < 0.001$ (two-tailed Student's *t*-test); NS, not significant. Scale bars, 50 nm. Data are mean \pm s.e.m. **f, g**, Effect of adding different CDC42 forms on the binding of cargo tails by coatomer; $n = 2$.

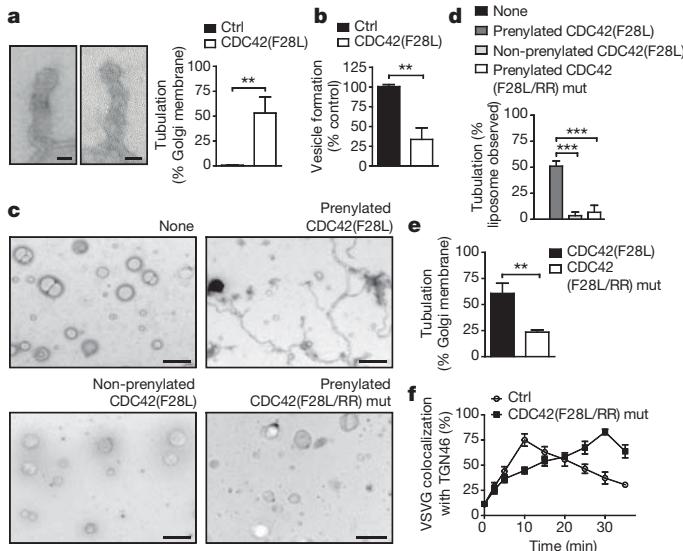


Figure 4 | CDC42 modulates carrier formation by COPI. **a**, Effect of active CDC42 on COPI tubule formation, with representative (out of 15) electron microscopy images and quantification shown; $n = 4$. Scale bars, 50 nm. ** $P < 0.01$ (two-tailed Student's t -test). **b**, Effect of active CDC42 on COPI vesicle formation; $n = 6$. ** $P < 0.01$ (two-tailed Student's t -test). **c**, **d**, Effect of mutations in CDC42 on its ability to induce liposome tubulation. **c**, Representative (out of 10) electron microscopy images. Scale bars, 500 nm. **d**, Quantification of data in **c**; $n = 5$. *** $P < 0.001$ (two-tailed Student's t -test). **e**, Effect of mutations in CDC42 on its ability to induce COPI tubule formation; $n = 4$. ** $P < 0.01$ (two-tailed Student's t -test). **f**, Effect of mutations in CDC42 on its ability to promote the intra-Golgi transport of VSVG; $n = 3$. Data are mean \pm s.e.m.

(Extended Data Fig. 7a), to be required for dimerization (Extended Data Fig. 7b). When the di-arginine residues were mutated, we found that active CDC42(Phe28Leu) can no longer tubulate liposomes (Fig. 4c, d). As further confirmation, we found that mutating the di-arginine residues also prevents CDC42 from promoting tubule formation in the COPI reconstitution system (Fig. 4e). Moreover, this mutation prevents CDC42 from promoting the intra-Golgi transport of VSVG (Fig. 4f and Extended Data Fig. 7c). We also considered that the di-arginine residues are closely juxtaposed to the di-lysine residues in CDC42 (Extended Data Fig. 7a). Thus, we assessed whether the di-arginine residues have a role in COPI cargo sorting. However, mutation of the di-arginine residues does not prevent activated CDC42(Phe28Leu) from interacting with coatomer (Extended Data Fig. 7d). Moreover, this mutation also does not alter the ability of active CDC42(Phe28Leu) to inhibit coatomer binding to the retrograde (KDEL) cargo (Extended Data Fig. 7e). In comparison, mutating the di-lysine residues in CDC42 affects this cargo binding by coatomer (Extended Data Fig. 7f). These binding results were also further confirmed in cells through FLIM (Extended Data Fig. 7g).

We next considered that coatomer has been identified as a target by which CDC42 promotes cellular transformation¹³. Led by this finding, we found that serum induces similar effects on bidirectional Golgi transport as activated CDC42 (Extended Data Fig. 8a, b). Moreover, epidermal growth factor (EGF) can replace serum in exerting these effects (Extended Data Fig. 8c, d). By contrast, neither serum (Extended Data Fig. 8e) nor EGF (Extended Data Fig. 8f) affects transport from the ER to the *cis*-Golgi. Mechanistically, serum acts through CDC42, as the enhanced anterograde transport of VSVG (Extended Data Fig. 8a) and the inhibited retrograde transport of VSVG-KDEL (Extended Data Fig. 8b) are prevented after CDC42 siRNA. This was also the case for EGF stimulation (Extended Data Fig. 8c, d).

We then considered that the SRC kinase has also been shown to promote anterograde Golgi transport¹⁹. Using siRNA targeting the

SRC gene, we found that serum can no longer modulate bidirectional Golgi transport (Extended Data Fig. 9a, b). We also found that pharmacological activation of SRC reproduces the effects of serum (Extended Data Fig. 9c, d). Moreover, *CDC42* siRNA prevents SRC activation from exerting these effects (Extended Data Fig. 9c, d). Thus, we concluded that a signalling cascade, which involves growth factors activating SRC and then CDC42, modulates bidirectional Golgi transport (Extended Data Fig. 10a). Notably, as siRNA against cPLA₂ prevents the ability of SRC activation to promote VSVG transport through the Golgi (Extended Data Fig. 10b), we further concluded that the signalling cascade targets COPI tubular transport in promoting anterograde Golgi transport.

In summary, we have advanced a new understanding of how Golgi transport is achieved and regulated. Our findings also address a continuing controversy about the role of COPI in anterograde Golgi transport. Although this debate has focused on the role of COPI vesicular transport²⁰, we provide a reconciling explanation by showing how COPI acts anterograde through tubular transport. Another notable finding is that CDC42 acts not just as a regulator of cellular events, but also as an effector, which in the case of COPI transport involves a direct ability to impart membrane curvature. Our results also shed new insight into how CDC42 targets coatomer for cellular transformation. Whereas CDC42 was previously thought to affect COPI transport in the retrograde direction¹³, our finding that CDC42 modulates bidirectional COPI transport to favour the anterograde direction suggests more readily how transformation is promoted, as the rapid growth and division of tumours is predicted to require enhanced secretion.

Online Content Methods, along with any additional Extended Data display items and Source Data, are available in the online version of the paper; references unique to these sections appear only in the online paper.

Received 13 August 2014; accepted 10 April 2015.

Published online 6 May 2015.

- Glick, B. S. & Nakano, A. Membrane traffic within the Golgi apparatus. *Annu. Rev. Cell Dev. Biol.* **25**, 113–132 (2009).
- Nakano, A. & Luini, A. Passage through the Golgi. *Curr. Opin. Cell Biol.* **22**, 471–478 (2010).
- Bonfanti, L. et al. Procollagen traverses the Golgi stack without leaving the lumen of cisternae: evidence for cisternal maturation. *Cell* **95**, 993–1003 (1998).
- Losev, E. et al. Golgi maturation visualized in living yeast. *Nature* **441**, 1002–1006 (2006).
- Matsuura-Tokita, K., Takeuchi, M., Ichihara, A., Mikuriya, K. & Nakano, A. Live imaging of yeast Golgi cisternal maturation. *Nature* **441**, 1007–1010 (2006).
- Patterson, G. H. et al. Transport through the Golgi apparatus by rapid partitioning within a two-phase membrane system. *Cell* **133**, 1055–1067 (2008).
- Trucco, A. et al. Secretory traffic triggers the formation of tubular continuities across Golgi sub-compartments. *Nature Cell Biol.* **6**, 1071–1081 (2004).
- San Pietro, E. et al. Group IV phospholipase A₂ controls the formation of inter-cisternal continuities involved in intra-Golgi transport. *PLoS Biol.* **7**, e1000194 (2009).
- Yang, J. S. et al. COPI acts in both vesicular and tubular transport. *Nature Cell Biol.* **13**, 996–1003 (2011).
- Hsu, V. W., Lee, S. Y. & Yang, J. S. The evolving understanding of COPI vesicle formation. *Nature Rev. Mol. Cell Biol.* **10**, 360–364 (2009).
- Cosson, P. & Letourneur, F. Coatomer interaction with di-lysine endoplasmic reticulum retention motifs. *Science* **263**, 1629–1631 (1994).
- Yang, J. S. et al. ARFGAP1 promotes the formation of COPI vesicles, suggesting function as a component of the coat. *J. Cell Biol.* **159**, 69–78 (2002).
- Wu, W. J., Erickson, J. W., Lin, R. & Cerione, R. A. The γ -subunit of the coatomer complex binds Cdc42 to mediate transformation. *Nature* **405**, 800–804 (2000).
- Yang, J. S. et al. A role for BARS at the fission step of COPI vesicle formation from Golgi membrane. *EMBO J.* **24**, 4133–4143 (2005).
- Yang, J. S. et al. A role for phosphatidic acid in COPI vesicle fission yields insights into Golgi maintenance. *Nature Cell Biol.* **10**, 1146–1153 (2008).
- Zimmerberg, J. & Kozlov, M. M. How proteins produce cellular membrane curvature. *Nature Rev. Mol. Cell Biol.* **7**, 9–19 (2006).
- McMahon, H. T. & Gallop, J. L. Membrane curvature and mechanisms of dynamic cell membrane remodelling. *Nature* **438**, 590–596 (2005).
- Zhang, B. & Zheng, Y. Negative regulation of Rho family GTPases Cdc42 and Rac2 by homodimer formation. *J. Biol. Chem.* **273**, 25728–25733 (1998).
- Pulvirenti, T. et al. A traffic-activated Golgi-based signalling circuit coordinates the secretory pathway. *Nature Cell Biol.* **10**, 912–922 (2008).
- Emr, S. et al. Journeys through the Golgi—taking stock in a new era. *J. Cell Biol.* **187**, 449–453 (2009).

Supplementary Information is available in the online version of the paper.

Acknowledgements We thank J. Li, M. Bai, X. Michelet and C. Alves for discussions, and M. Ericsson for electron microscopy technical advice. This work was funded by grants from the National Institutes of Health to V.W.H. (R01GM058615), R.J.S. (R01AI068871, R01AR065538 and 1S10RR027931-01), A.B.S. (K01DK089145), and also by the Basic Science Research Program of the National Research Foundation of Korea to S.-Y.P. (2014R1A6A3A03056673).

Author Contributions S.-Y.P., J.-S.Y. and A.B.S. performed the experiments. V.W.H., S.-Y.P., R.J.S. and A.B.S. designed the experiments and wrote the paper.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to V.W.H. (vhsu@research.bwh.harvard.edu).

METHODS

Chemicals, lipids, peptides and cells. The SRC family activator (Santa Cruz, SC3052) has been described²². GDP and GTP were purchased (Sigma-Aldrich). DOPC (dioleoylphosphatidylcholine), DOPE (dioleoylphosphatidylethanolamine), DOPS (dioleoylphosphatidylserine), phosphatidic acid and cholesterol were also purchased (Avanti Polar Lipid). Peptides encoding the WBP1 tail (wild-type, KKLETFKKT; mutant, KKLETSSTN) were synthesized (GenScript). Culturing of HeLa cells has been described^{9,14,15}.

Antibodies. Mouse anti- β subunit of coatomer, anti- ζ subunit of coatomer, mouse anti-Myc (9E10), sheep anti-TGN46, rabbit anti-giantin antibodies have been described previously²¹. Rabbit anti-Myc antibody (2278S), rabbit anti-CDC42 antibody (2462S), and mouse anti- β -actin antibody (3700P) were obtained (Cell Signaling). All secondary antibodies were also obtained (Jackson ImmunoResearch Laboratory).

Purified and recombinant proteins. Preparation of coatomer, ARF1, ARFGAP1, BARS, cPLA_{2 α} and Golgi membrane have been described^{9,12,14}. To generate the prenylated form of human CDC42, pFASTBAC THb (Invitrogen) was used to generate baculovirus using Bac-to-Bac Baculovirus Expression system (Invitrogen). Infected SF9 cells were resuspended in hypotonic buffer (20 mM sodium borate, pH 10.2, 5 mM MgCl₂, and protease inhibitor) followed by lysis using dounce homogenizer. Membrane-bound CDC42 was collected by 150,000g centrifugation for 30 min. The pellet was then resuspended with TBS-M (50 mM Tris (pH 7.2), 150 mM NaCl, 5 mM MgCl₂ and protease inhibitor), followed by centrifugation at 150,000g for 30 min. Prenylated CDC42 was solubilized from membranes using TBS-M containing 1% Triton X-100 for 1 h, followed by centrifugation at 9,000g for 20 min. The histidine-tagged CDC42 was then purified from the supernatant fraction using a nickel column. Non-prenylated CDC42 was generated using bacterial expression. The last three amino acids were deleted to mimic the protein sequence of the prenylated form (see Extended Data Fig. 7a). Human CDC42 in pET-15b (Novagen) was expressed in bacteria (BL21) with 100 μ M of isopropyl- β -D-thiogalactopyranoside (IPTG) induction at 37 °C for 2 h. The His-tagged CDC42 was isolated with a nickel column, and then released with elution buffer (50 mM Tris (pH 7.2), 500 mM NaCl, 5 mM MgCl₂, 0.1% CHAPS and 250 mM imidazole). Recombinant forms of RAC1, RhoA and RAB1A, RAB2A, RAB6A and RAB33B were generated similarly.

Plasmids. Cytoplasmic domain of cargoes fused to glutathione S-transferase (GST) was generated using *in vitro* annealing. Paired oligonucleotides were annealed in buffer (30 mM HEPES-KOH, pH 7.4, 150 mM KCl, 2 mM MgCl₂) at 90 °C for 1 min followed by cooling down. Annealed duplexes were inserted into pGEX-4T-3 (GE Healthcare). GST-fused cytoplasmic peptides of LDLR were generated using PCR. VSVG-LDLR was generated by fusing the cytoplasmic tail of LDLR to the luminal and transmembrane domains of VSVG-ts045. Mutations were generated using PCR or QuikChange Site-Directed-Mutagenesis (Stratagene). VSVG-ts045 and VSVG-ts045-KDEL were described previously²¹. CDC42(Phe28Leu) mutant, and the di-lysine mutant of CDC42(Phe28Leu) in pFASTBAC THb and PKH3 were gifts from R. A. Cerione. Di-arginine mutant of CDC42 was generated using PCR with paired oligonucleotides. Additional plasmids were obtained and then subcloned into pET-15b to generate recombinant proteins: RAC1 and RhoA (D. Manor), RAB1A, RAB2A and RAB6A (D. Lambright), and RAB33B (M. Fukuda). See Supplementary Information for details on the oligonucleotides used to generate chimaeric constructs and point mutations.

Transfections and siRNA. Transfection of DNA plasmids was performed using FuGene6 (Roche). Transfection of siRNA was carried out using Lipofectamine RNAiMAX (Invitrogen). The siRNA sequence against cPLA_{2 α} (targeting the PLA2G4A gene) has been described previously⁹. The siRNA sequences used to target human CDC42 (5'-GATAACTCACCACTGTCCA-3') and human SRC (5'-TGTCGGAGGCTTCAACTCCT-3') were obtained (Dharmacon). Point mutations in CDC42 to render siRNA resistant were generated using oligonucleotides (5'-GAAAGTGGGTGCCTGAGATAACTCATCACTGTCAAAGAC TCTTTCTGCTTGGG-3' and 5'-CCCAACAAGCAAGAAAGGAGTC TTTGGACAGTGTGAGTTATCTCAGGCACCCACTTTC-3').

In vitro binding assays. GST fusions were expressed in bacteria (BL21 cell) with IPTG induction. After lysis in buffer (20 mM Tris, pH 8.0, 100 mM NaCl, 1 mM EDTA, 1% Triton X-100, 1 mg ml⁻¹ of lysozyme and protease inhibitor), GST peptide was incubated with glutathione Sepharose beads (GE Healthcare). GST fusions on beads were then incubated with purified coatomer (2.5 nM) at 4 °C for 1 h in incubation buffer (25 mM HEPES, pH 7.2, 50 mM KCl, 2.5 mM magnesium acetate and 0.5% NP-40). For competition experiments with peptides, GST fusions on beads were co-incubated with coatomer and peptide at 4 °C for 1 h in incubation buffer. Beads were then rinsed twice with incubation buffer and then analysed by SDS-PAGE followed by immunoblotting or Coomassie staining. For competition experiments with small GTPases, GTPases were loaded with either GTP or

GDP and then incubated with cargo (as GST fusions on beads) and coatomer. Quantification was performed by analysing the level of coatomer and normalizing for the level of cargo on beads.

In vivo transport assays. Anterograde transport of VSVG and retrograde transport of VSVG-KDEL have been described previously^{9,21}. In brief, HeLa cells were transfected with VSVG-ts045 and VSVG-ts045-KDEL to track anterograde and retrograde transport, respectively. To track transport from the ER to the Golgi, cells were incubated at 40 °C for 2 h to accumulate a synchronized pool of VSVG at the ER. Cells were then shifted to incubation at 32 °C to allow transport from the ER to proceed. To examine intra-Golgi transport, cells were incubated at 15 °C for 2 h to accumulate a synchronized pool of VSVG at the pre-Golgi region, followed by a shift to 32 °C to allow transport through the Golgi. To study retrograde transport, cells expressing VSVG-ts045-KDEL were incubated at 32 °C for 2 h. Cells were then incubated at 40 °C. This prevents further anterograde transport from the ER, and thereby allowing one round of retrograde transport from the Golgi to the ER to be assessed. In stimulation experiments, cells were starved using DMEM supplemented with 0.1% BSA for 2 h followed by stimulation with serum or a SRC activator phosphopeptide²² (10 μ M) for 1 h. Transport was quantified by measuring (Metamorph) the colocalization of cargoes with organelle markers over time.

FLIM. Interactions were monitored using time-correlated single-photon counting fluorescence lifetime image microscopy analysis (TCSPC-FLIM), which has been described previously^{23,24}. In brief, the different forms of VSVG were detected through the Myc tag appended at the C terminus using an anti-Myc antibody. This antibody was then detected with a secondary antibody conjugated to Alexa Fluor 594 (acceptor fluorophore). Coatomer was detected with the anti-coatomer antibody. This antibody was then detected with a secondary antibody conjugated to Alexa Fluor 488 (donor fluorophore). The baseline lifetimes of the donor fluorophore were calculated by single-exponential decay fitting of fluorescence emission in the absence of the acceptor fluorophore. For samples that involved staining for both donor and acceptor, lifetimes were fitted to a bi-exponential decay with lifetime of one component fixed to the donor-only lifetime. Three variables for FLIM were determined: (1) lifetime for the interacting fraction, τ_1 , (2) lifetime for the non-interacting fraction, τ_2 , and (3) the percentage of interacting molecules, a_1 (%). Because VSVG is transported from the ER to the Golgi, we used this mobility to define key physiological parameters in quantifying the interaction between VSVG and coatomer (based on the τ_1 values). The maximal interaction was defined by quantifying VSVG interacting with coatomer at the Golgi, and minimal interaction was defined by performing quantitation when VSVG is at the ER. Interactions of the different forms of VSVG with coatomer were then expressed as the fraction of the maximal interaction (coatomer interacting with wild-type VSVG at the Golgi).

COPI reconstitution system. The two-stage incubation system was performed as described previously^{9,14}. In brief, Golgi membrane (250 μ g) was washed with 100 μ l of 3 M KCl for 5 min, and then diluted with 1 ml of reaction buffer (25 mM HEPES-KOH, pH 7.2, 50 mM KCl, 2.5 mM magnesium acetate, 1 mg ml⁻¹ soybean trypsin inhibitor, 1 mg ml⁻¹ BSA and 200 mM sucrose). In the first stage, the washed membrane was incubated with ARF1, coatomer and GTP at 37 °C for 10 min. Reaction was stopped on the ice water for 5 min, followed by centrifugation at 12,000g for 20 min. In the second stage, the pellet was resuspended with reaction buffer, and then incubated with ARFGAP1 and BARS for 20 min to generate COPI vesicles, or additionally with CDC42 and/or cPLA_{2 α} to generate COPI tubules.

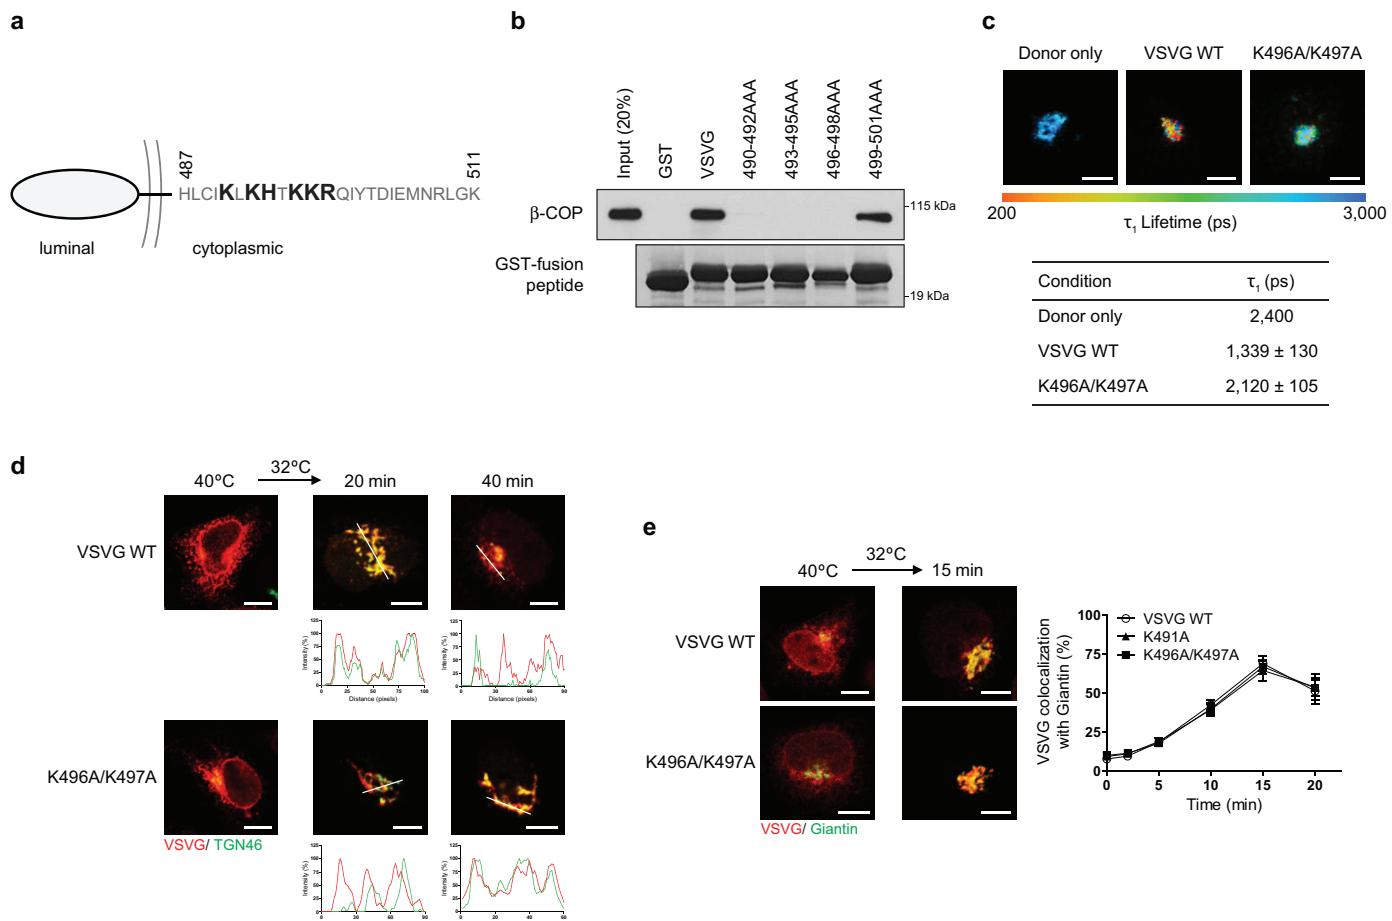
Liposome tubulation assay. Liposomes were generated using a mixture of pure lipids (Avanti Polar Lipids): DOPC (33%), DOPS (5%), DOPE (15%), phosphatidic acid (30%) and cholesterol (18%). Lipids (500 μ g) in glass tubes were evaporated under nitrogen gas. Lipids were resuspended with assay buffer (25 mM HEPES, pH 7.2, 50 mM KCl, 2.5 mM magnesium acetate and 200 mM sucrose) and hydrated overnight. Liposomes were then passed through 400-nm filter membrane (Whatman) in a mini-extruder (Avanti Polar Lipids). Recombinant CDC42 was then incubated with liposomes at room temperature for 30 min, followed by electron microscopy analysis.

Electron microscopy analysis. Electron microscopy analysis of Golgi tubules has been described previously⁹. In brief, membrane samples were loaded on electron microscopy grids and then fixed with 2% paraformaldehyde (PFA) in PBS for 10 min. After blocking with 1% BSA in PBS for 10 min, grids were rinsed with water, followed by uranyl acetate staining. Golgi tubule was detected using JEOL 1200EX transmission electron microscope. To detect cargoes in the reconstituted tubules, immunogold labelling was performed, which involved incubation with mouse anti-Myc antibody (9E10) for 1 h to detect the Myc-tagged form of either VSVG or VSVG-KDEL. After rinsing with PBS, the sample was incubated with a rabbit anti-mouse antibody for 30 min. After an additional rinse

with PBS, the sample was further incubated with protein A (conjugated with 10 nm gold particles). Grids were then washed with water, followed by uranyl acetate staining, and then examination by electron microscopy. Fifty electron microscopy meshes were examined for quantification in each condition of an experiment.

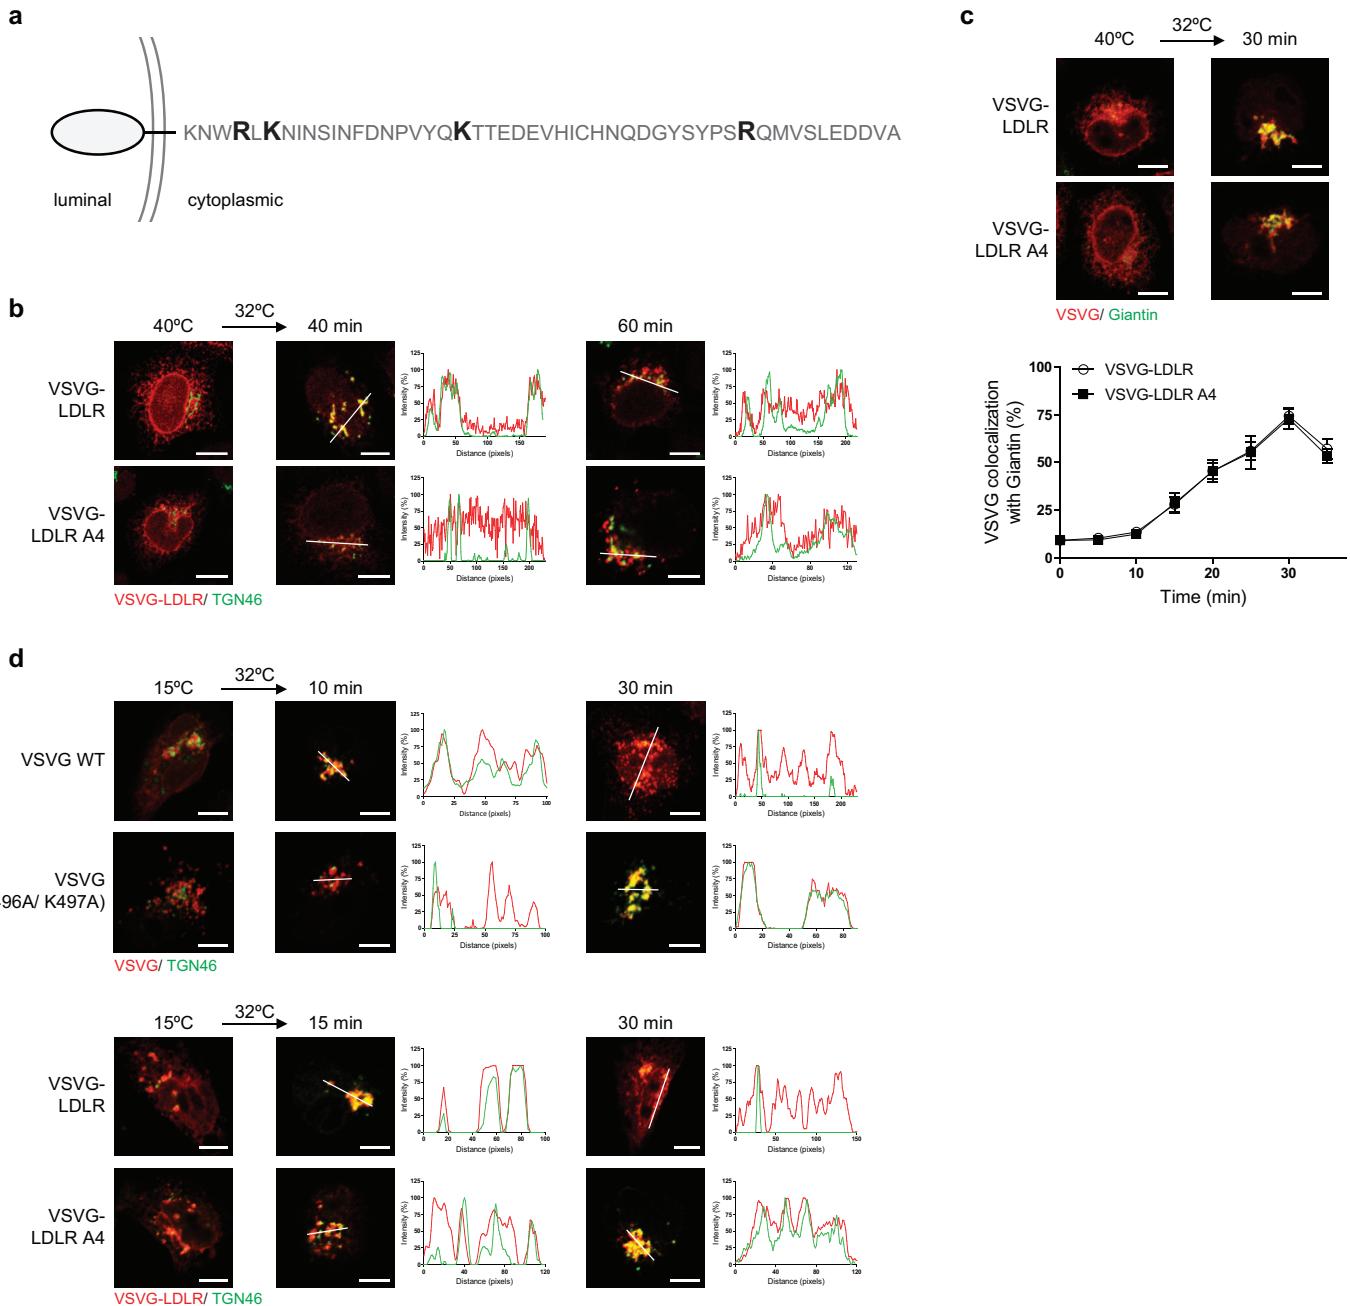
Gel filtration chromatography. Dimerization of CDC42 was determined using Superdex 200 10/300 GL (GE Healthcare), essentially as previously described¹⁸. Recombinant CDC42 was applied at a flow rate of 0.5 ml min⁻¹ and monitored by ultraviolet detector.

21. Bai, M. *et al.* ARFGAP1 promotes AP-2 dependent endocytosis. *Nature Cell Biol.* **13**, 559–567 (2011).
22. Lu, B. *et al.* Peptide neurotransmitters activate a cation channel complex of NALCN and UNC-80. *Nature* **457**, 741–744 (2009).
23. Bair, A. M., Turman, M. V., Vaine, C. A., Panettieri, R. A. Jr & Soberman, R. J. The nuclear membrane leukotriene synthetic complex is a signal integrator and transducer. *Mol. Biol. Cell* **23**, 4456–4464 (2012).
24. Azcutia, V. *et al.* CD47 plays a critical role in T-cell recruitment by regulation of LFA-1 and VLA-4 integrin adhesive functions. *Mol. Biol. Cell* **24**, 3358–3368 (2013).



Extended Data Figure 1 | Characterizing VSVG transport. **a**, Cytoplasmic sequence of VSVG with residues crucial for binding by coatomer highlighted. **b**, Effect of point mutations in the VSVG tail on the *in vitro* binding of this tail by coatomer, $n = 5$. **c**, FLIM assessing interaction between VSVG and coatomer. Representative images that are pseudo-coloured based on τ_1 values are shown (out of 10). Quantification is also shown, $n = 3$. **d**, Effect of

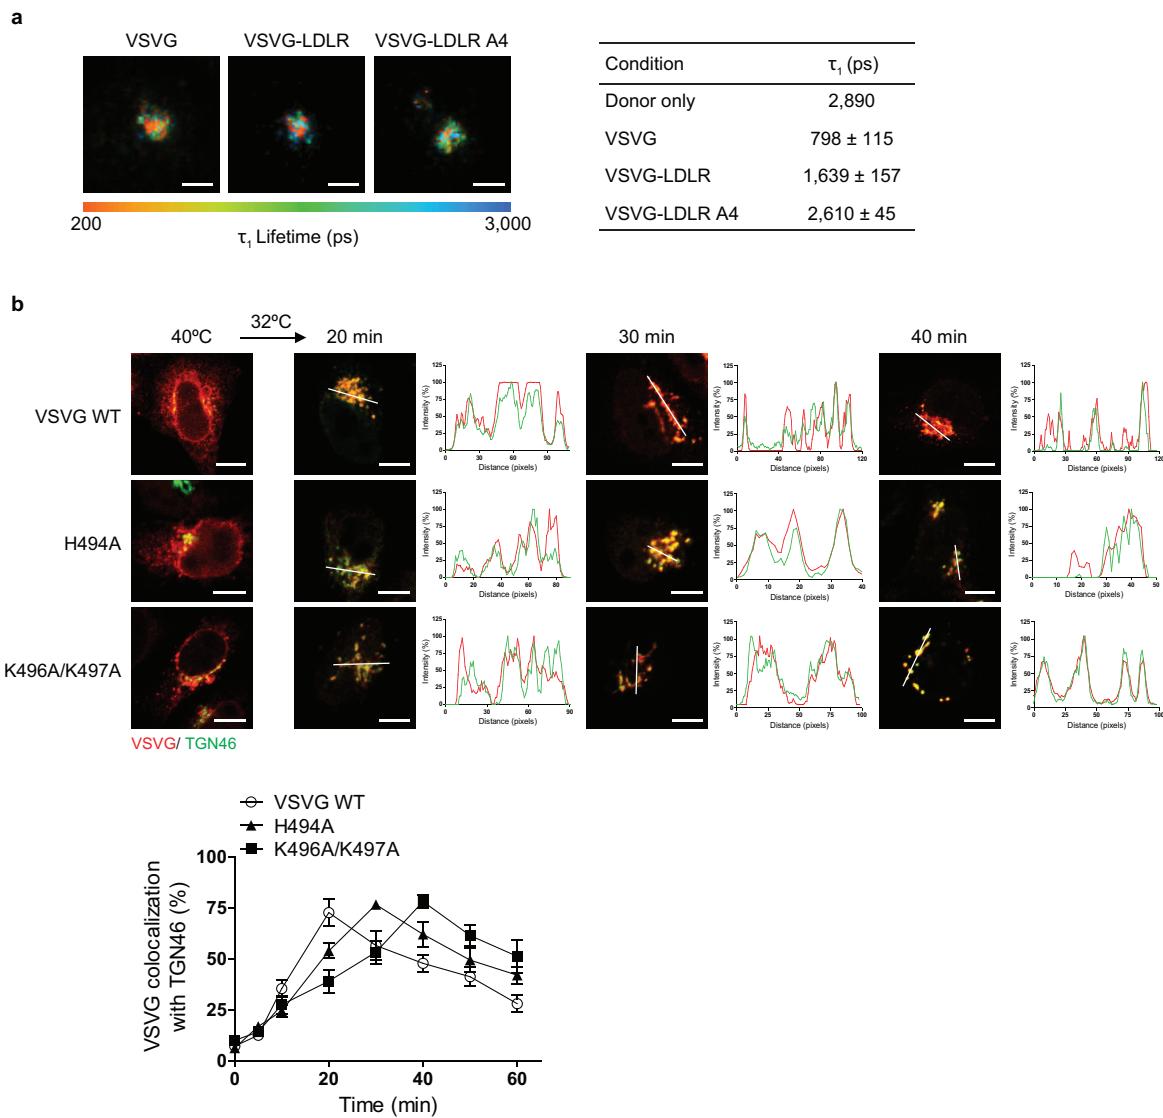
mutations in the VSVG tail on its transport from the ER to the *trans*-Golgi. Colocalization is exemplified by line scanning across representative images (out of 5). e, Effect of mutations in the VSVG tail on its transport from the ER to the *cis*-Golgi. Representative images of colocalization are shown (out of 5). Quantification is also shown, $n = 3$. Data are mean \pm s.e.m. Scale bars, 5 μm .



Extended Data Figure 2 | Characterizing VSVG–LDLR transport.

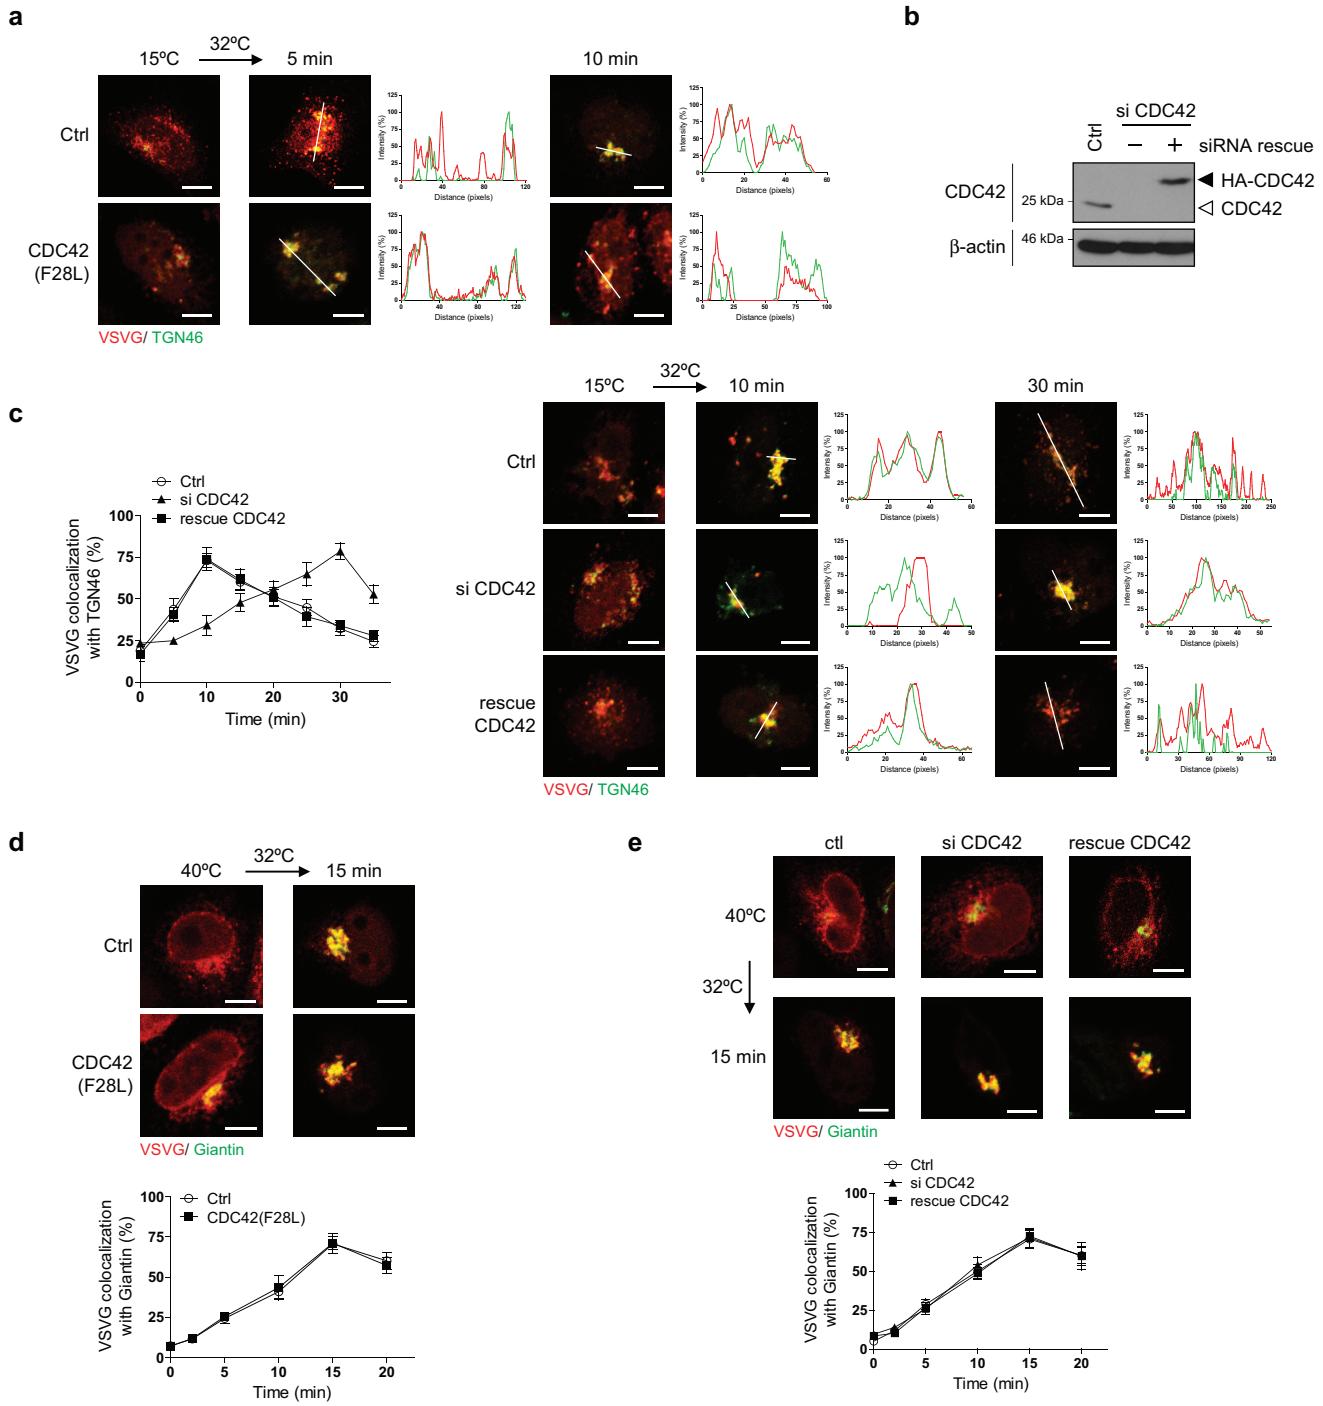
a, Cytoplasmic sequence of LDLR with basic residues critical for binding by coatomer highlighted. **b**, Effect of mutations in the LDLR tail on the transport of VSVG–LDLR from the ER to the *trans*-Golgi. Colocalization is exemplified by line scanning across representative images (out of 5). **c**, Effect of mutations in

the LDLR tail on the transport of VSVG–LDLR from the ER to the *cis*-Golgi. Representative images of colocalization are shown (out of 5). Quantification is also shown, $n = 4$, mean \pm s.e.m. **d**, Effect of mutations in cargo tails on the intra-Golgi transport of various cargoes. Colocalization is exemplified by line scanning across representative images (out of 5). Scale bars, 5 μm .



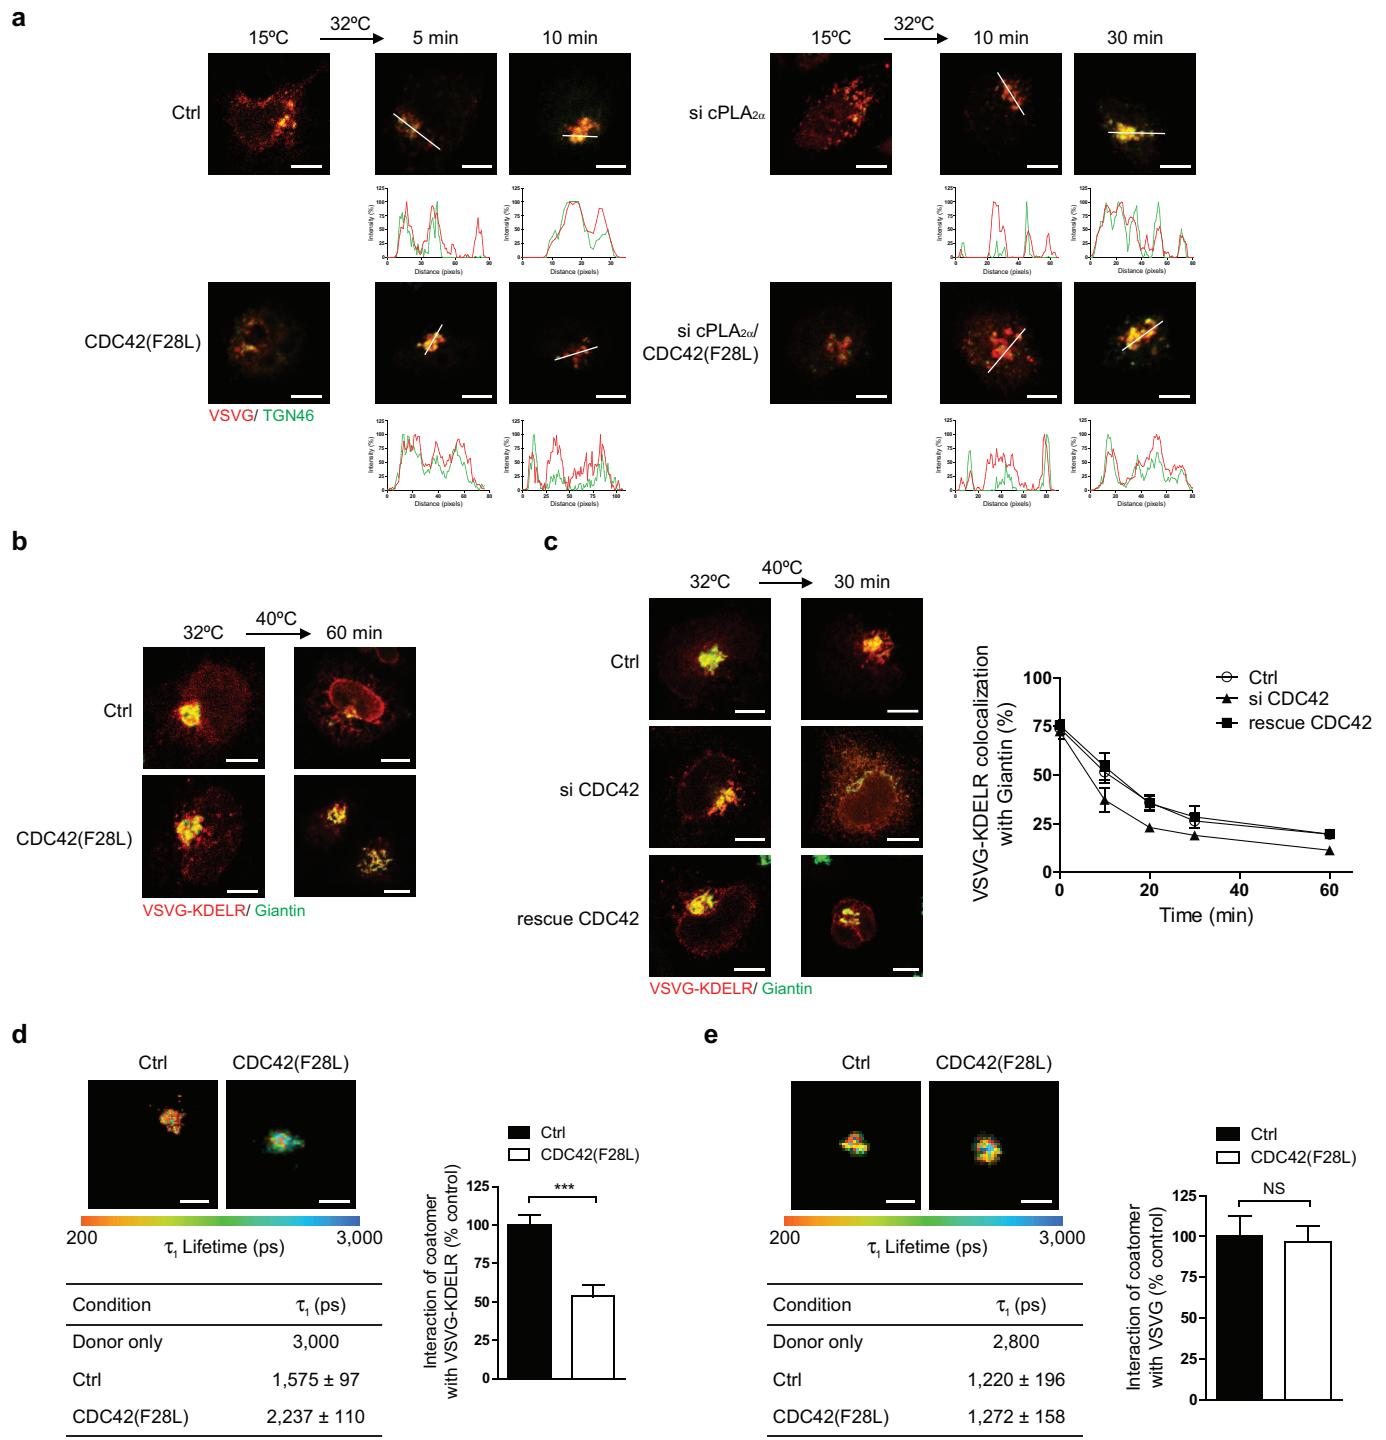
Extended Data Figure 3 | Further characterizing cargo transport. **a**, FLIM assessing interaction between different forms of VSVG and coatomer. Representative images that are pseudo-coloured based on τ_1 values are shown (out of 10). Quantification is also shown, $n = 3$. **b**, Effect of mutations in the

VSVG tail on its transport from the ER to the *trans*-Golgi. Colocalization is exemplified by line scanning across representative images (out of 5). Quantification is also shown, $n = 3$. Data are mean ± s.e.m. Scale bars, 5 μm .



Extended Data Figure 4 | Characterizing how CDC42 affects VSVG transport. **a**, Effect of expressing active CDC42 on the intra-Golgi transport of VSVG. Colocalization is exemplified by line scanning across representative images (out of 5). **b**, Immunoblotting of whole-cell lysates to assess efficiency of siRNA treatment, $n = 3$. **c**, Effect of siRNA against CDC42 on the intra-Golgi transport of VSVG. Colocalization is exemplified by line scanning across

representative images (out of 5). Quantification is also shown, $n = 4$. **d**, Effect of expressing active CDC42 on VSVG transport from the ER to the *cis*-Golgi. Representative images of colocalization are shown (out of 5). Quantification is also shown, $n = 5$. **e**, Effect of siRNA against CDC42 on VSVG transport from the ER to the *cis*-Golgi. Representative images of colocalization are shown (out of 5). Quantification is also shown, $n = 4$. Data are mean \pm s.e.m. Scale bars, 5 μ m.

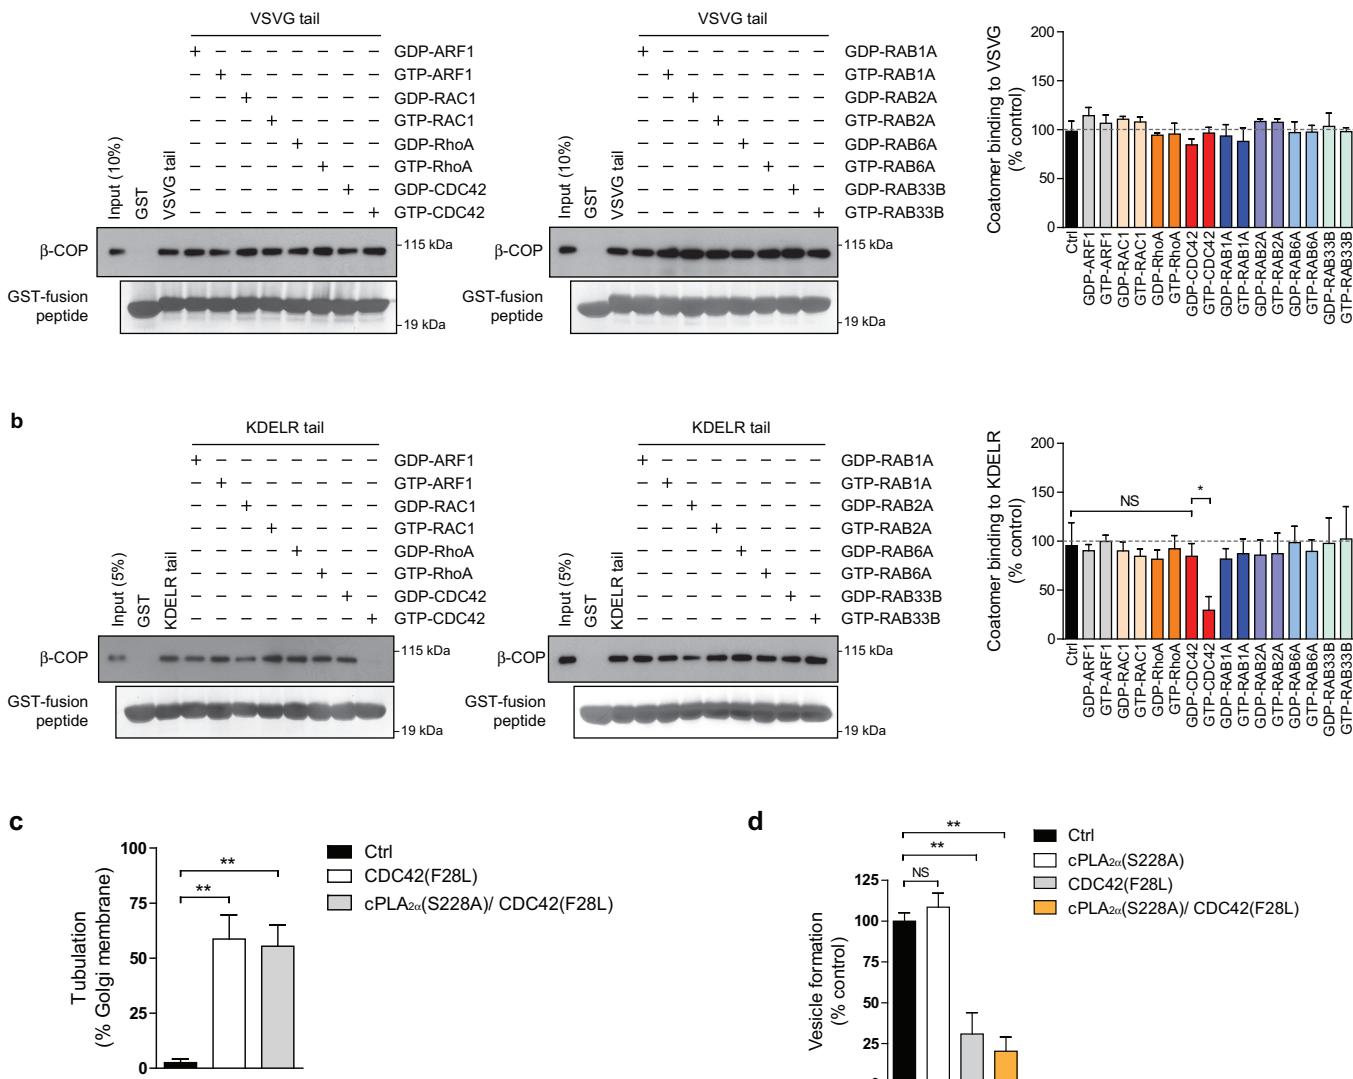


Extended Data Figure 5 | Further characterizing the effects of CDC42.

a, Effect of various conditions on the intra-Golgi transport of VSVG. Colocalization is exemplified by line scanning across representative images (out of 5). **b**, Effect of expressing active CDC42 on the transport of VSVG-KDELr. Representative images of colocalization is shown (out of 5). **c**, Effect of siRNA against CDC42 on the transport of VSVG-KDELr. Representative images of colocalization are shown (out of 5). Quantification is also shown, $n = 3$.

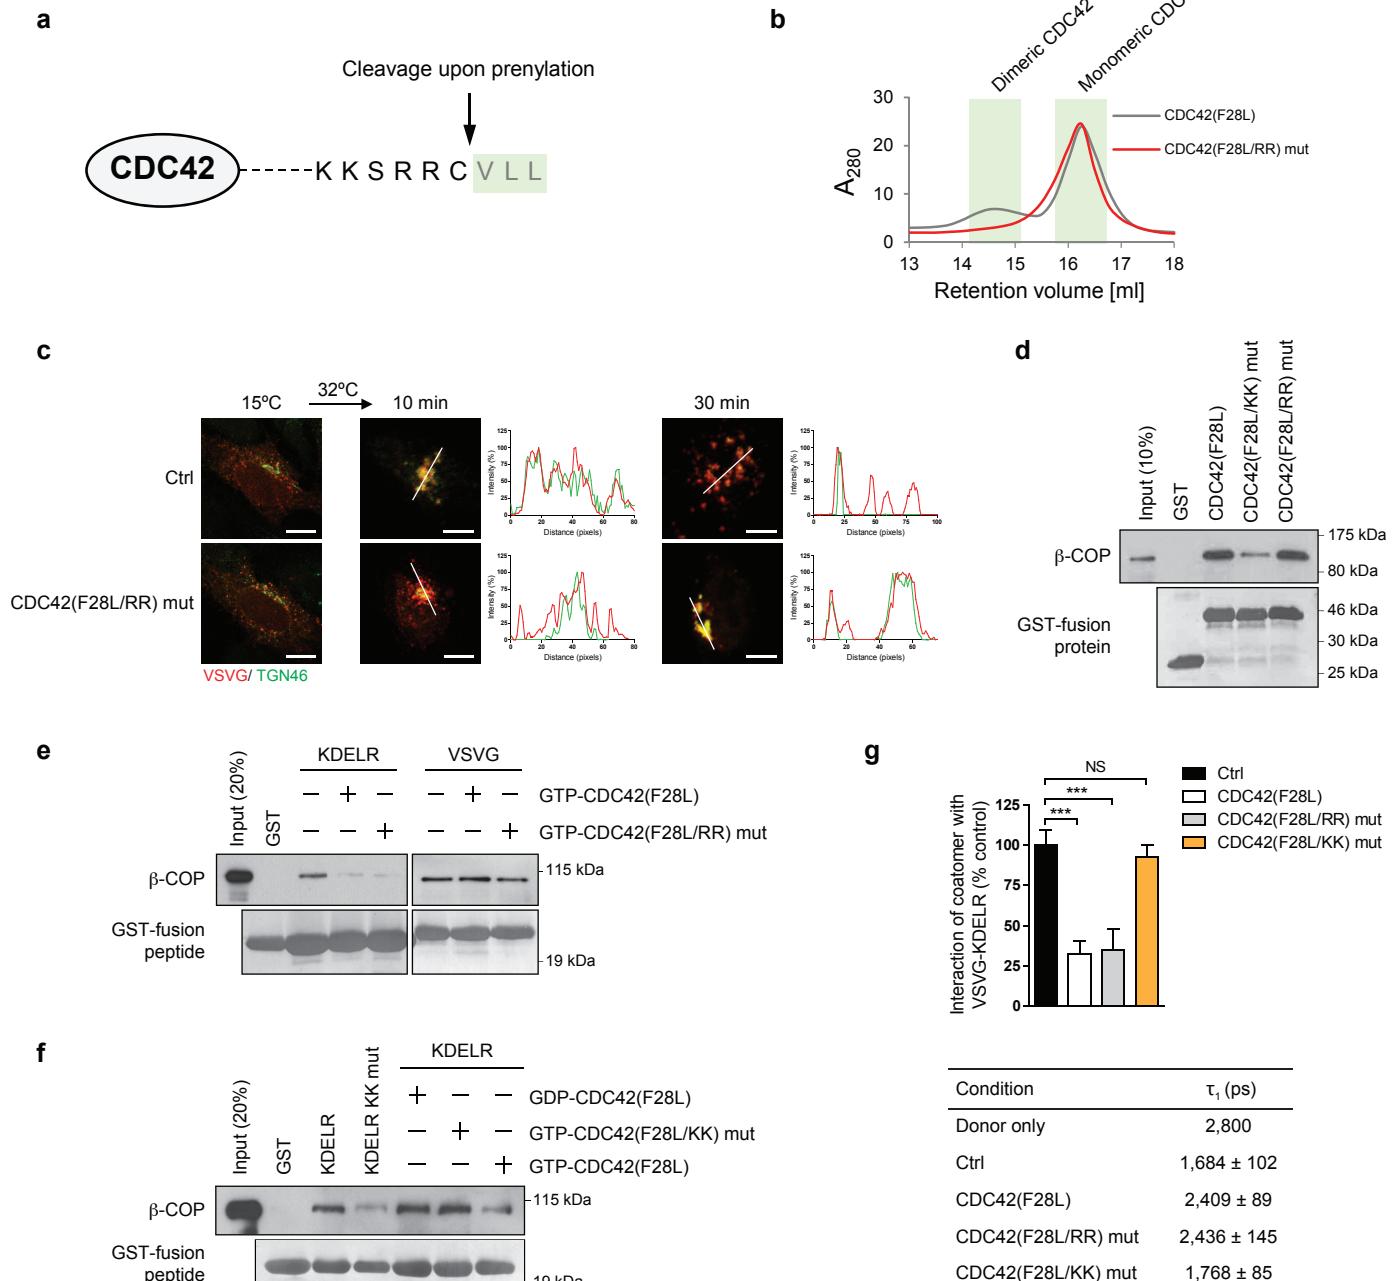
d, e, FLIM assessing the effect of active CDC42 on the interaction between different VSVG forms and coatomer. Representative images that are pseudo-coloured based on τ_1 values are shown (out of 10). Quantification is also shown as table and as graph, $n = 3$. *** $P < 0.001$ (two-tailed Mann–Whitney test). Data are mean \pm s.e.m. Scale bars, 5 μm .

a



Extended Data Figure 6 | Characterizing how CDC42 affects COPI cargo sorting and carrier formation. a, b, Effect of different small GTPases on the *in vitro* binding of cargo tails by coatomer. Representative blots are shown (out of 3). Quantification is also shown, $n = 3$, mean \pm s.e.m. * $P < 0.05$ (two-tailed

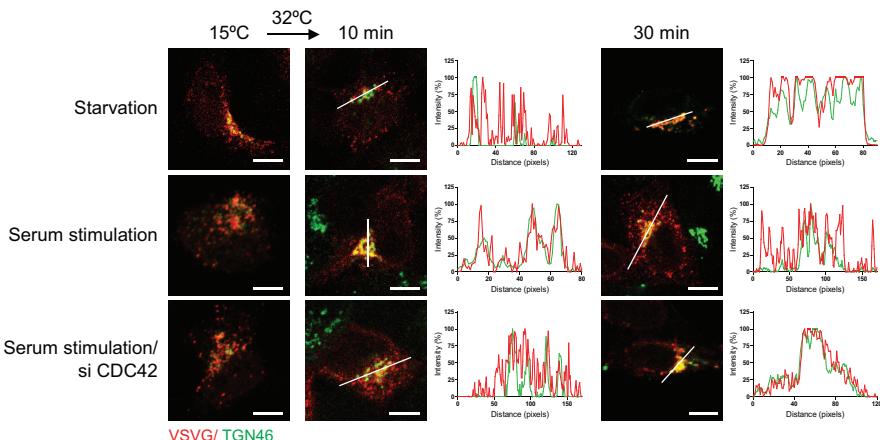
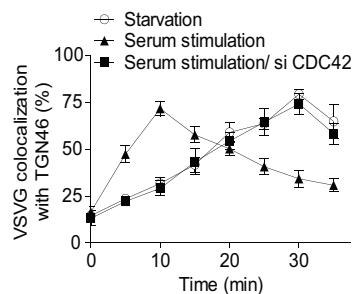
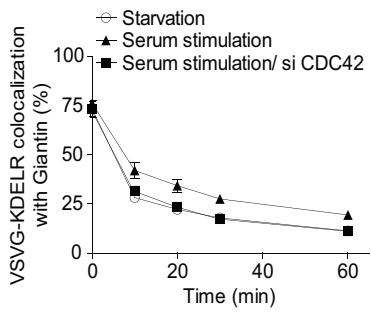
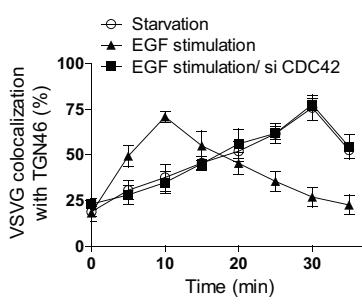
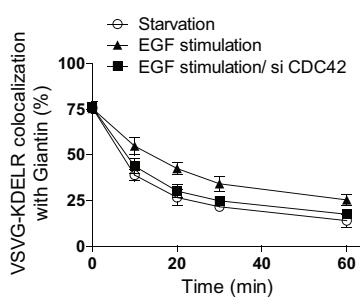
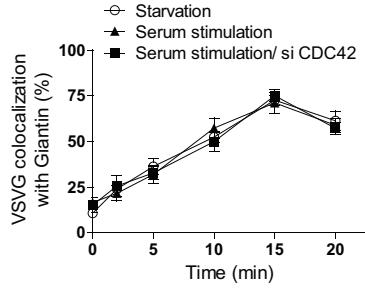
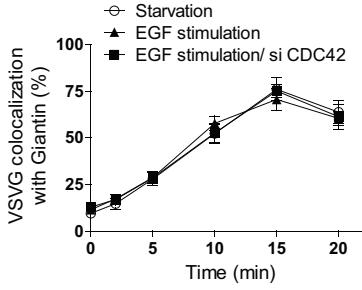
Student's *t*-test). **c, d**, Effect of different active forms of CDC42 on tubule and vesicle formation in the COPI reconstitution system, $n = 4$, mean \pm s.e.m.
 $^{**}P < 0.01$ (two-tailed Student's *t*-test).



Extended Data Figure 7 | Delineating the role of di-arginine residues in CDC42. **a**, Amino acid sequence at the C terminus of CDC42. The last three residues are cleaved after prenylation. **b**, The di-arginine residues in CDC42 are required for dimerization. Representative result from gel filtration (out of 2) is shown. **c**, Effect of different conditions on the intra-Golgi transport of VSVG. Colocalization is exemplified by line scanning across representative images (out of 5). Scale bar, 5 μ m. **d**, Effect of mutations in CDC42 on its ability to bind

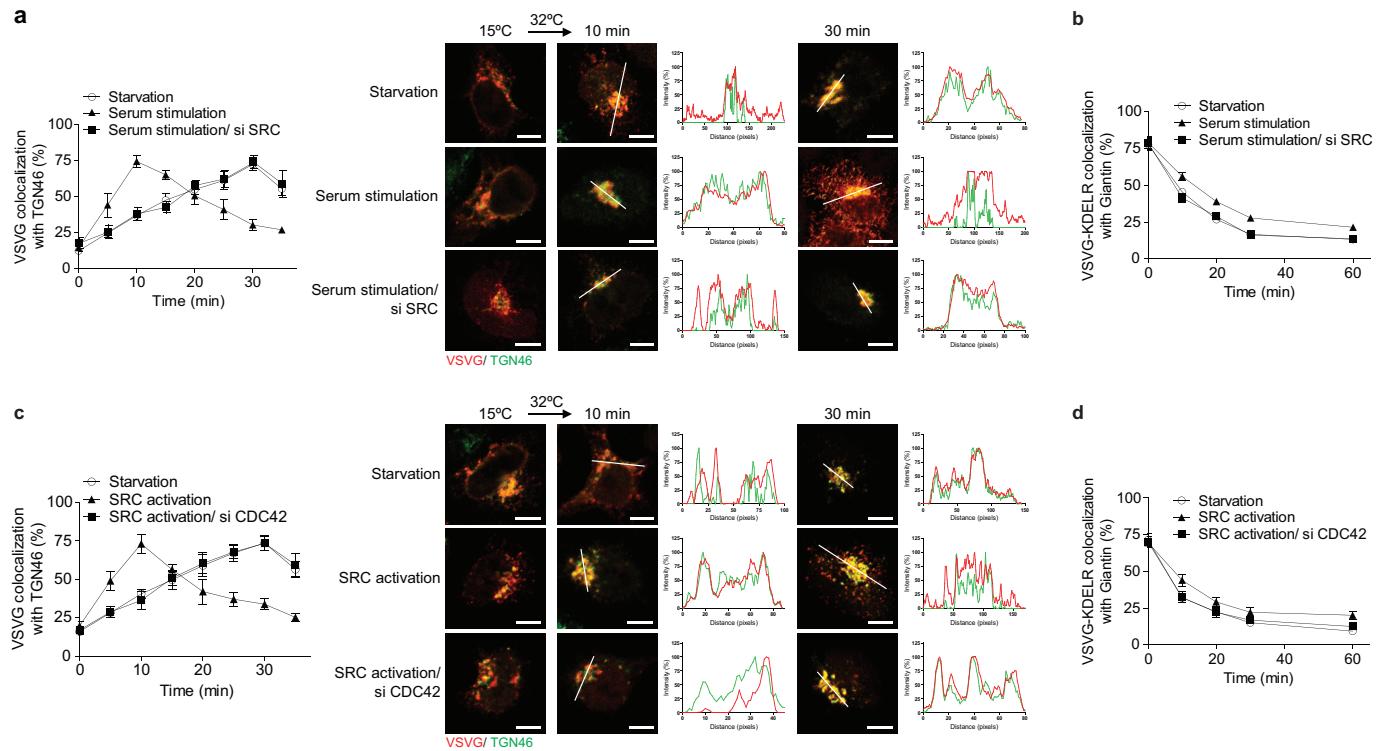
coatomer *in vitro*. Representative blot (out of 2) is shown. **e**, Effect of mutations in CDC42 on its ability to compete with cargo tails for binding to coatomer *in vitro*. Representative blot (out of 3) is shown. **f**, Effect of different forms of CDC42 to compete with the KDEL tail for binding to coatomer *in vitro*. Representative blot (out of 2) is shown. **g**, FLIM examining VSVG-KDEL interacting with coatomer, $n = 3$. Data are mean \pm s.e.m. *** $P < 0.001$ (two-tailed Mann–Whitney test).

Condition	τ_1 (ps)
Donor only	2,800
Ctrl	$1,684 \pm 102$
CDC42(F28L)	$2,409 \pm 89$
CDC42(F28L/RR) mut	$2,436 \pm 145$
CDC42(F28L/KK) mut	$1,768 \pm 85$

a**b****c****d****e****f**

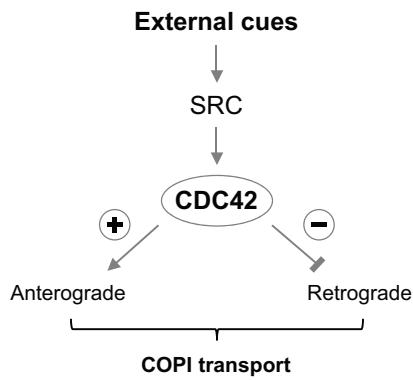
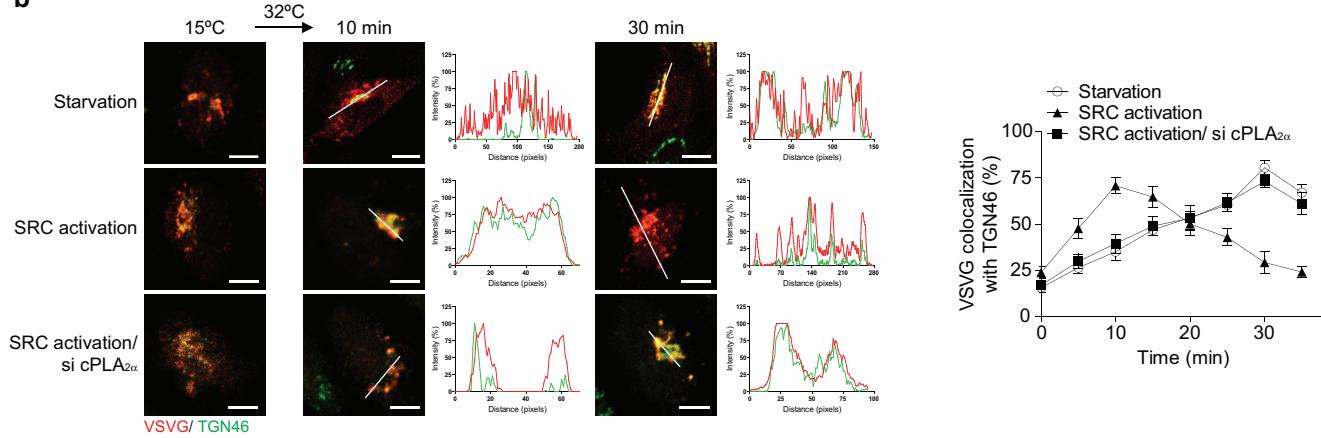
Extended Data Figure 8 | Delineating how external stimuli regulate bidirectional Golgi transport. **a**, Effect of different conditions on the intra-Golgi transport of VSVG. Colocalization is exemplified by line scanning across representative images (out of 5). Scale bar, 5 μ m. Quantification is also shown,

$n = 3$. **b**, Effect of serum on the transport of VSVG–KDEL R, $n = 3$. **c**, Effect of EGF on the intra-Golgi transport of VSVG, $n = 3$. **d**, Effect of EGF on the transport of VSVG–KDEL R, $n = 3$. **e,f**, Effect of different conditions on VSVG transport from the ER to the cis-Golgi, $n = 3$. Data are mean \pm s.e.m.



Extended Data Figure 9 | Delineating how SRC regulates bidirectional Golgi transport. **a, c,** Effect of different conditions on the intra-Golgi transport of VSVG. Colocalization is exemplified by line scanning across representative

images (out of 5). Scale bar, 5 μ m. Quantification is also shown, $n = 3$ experiments. **b, d,** Effect of different conditions on the transport of VSVG-KDEL, $n = 3$. Data are mean \pm s.e.m.

a**b**

Extended Data Figure 10 | Further characterizing how bidirectional COPI transport at the Golgi is regulated. **a**, Summarizing how external stimuli regulates bidirectional COPI transport through a signalling cascade. **b**, Effect of

SRC activation on the intra-Golgi transport of VSVG. Colocalization is exemplified by line scanning across representative images (out of 5). Scale bar, 5 μ m. Quantification is also shown, $n = 3$. Data are mean \pm s.e.m.

Defining fundamental steps in the assembly of the *Drosophila* RNAi enzyme complex

Shintaro Iwasaki^{1†*}, Hiroshi M. Sasaki^{1*}, Yuriko Sakaguchi², Tsutomu Suzuki², Hisashi Tadakuma³ & Yukihide Tomari^{1,3}

Small RNAs such as small interfering RNAs (siRNAs) and microRNAs (miRNAs) silence the expression of their complementary target messenger RNAs^{1,2} via the formation of effector RNA-induced silencing complexes (RISCs), which contain Argonaute (Ago) family proteins at their core. Although loading of siRNA duplexes into *Drosophila* Ago2 requires the Dicer-2-R2D2 heterodimer^{3–5} and the Hsc70/Hsp90 (Hsp90 also known as Hsp83) chaperone machinery^{6–8}, the details of RISC assembly remain unclear. Here we reconstitute RISC assembly using only Ago2, Dicer-2, R2D2, Hsc70, Hsp90, Hop, Droj2 (an Hsp40 homologue) and p23. By following the assembly of single RISC molecules, we find that, in the absence of the chaperone machinery, an siRNA bound to Dicer-2-R2D2 associates with Ago2 only transiently. The chaperone machinery extends the dwell time of the Dicer-2-R2D2-siRNA complex on Ago2, in a manner dependent on recognition of the 5'-phosphate on the siRNA guide strand. We propose that the chaperone machinery supports a productive state of Ago2, allowing it to load siRNA duplexes from Dicer-2-R2D2 and thereby assemble RISC.

Because both miRNAs and siRNAs are generated as double-stranded RNAs—miRNA/miRNA* duplexes and siRNA duplexes—but mature RISCs contain only single-stranded guide RNAs, RISC assembly involves at least two steps: loading of a small RNA duplex into Ago (forming pre-RISC) and ejection of the passenger strand from Ago (forming mature RISC)^{1,2}. Even when initiated by ‘pre-diced’ siRNA duplexes rather than long double-stranded RNAs, assembly of *Drosophila* Ago2-RISC requires Dicer-2, an RNase III enzyme that processes long double-stranded RNAs into siRNA duplexes, and its partner protein R2D2 (refs 3–5). In contrast, Dicer or its partner proteins are not essential for assembly of fly Ago1-RISC⁹ and mouse and human AGO2-RISC^{10–14}. Regardless of Dicer dependence, loading of small RNA duplexes into Ago requires ATP^{4–7,9,15} and is blocked by Hsp70 or Hsp90 inhibitors^{6–8,16}. On the other hand, subsequent ejection of the passenger strand from Ago proceeds without ATP or chaperones^{7,9,15}, but is accelerated by cleavage of the passenger strand by Ago^{17–19} and the endonuclease C3PO^{12,20}.

The Hsc70 and Hsp90 chaperones generally act together with other multiple proteins, which as a large complex help the client proteins achieve an active conformation²¹. For example, maturation of mammalian steroid hormone receptors to an active ligand-binding form requires HSC70, HSP90, HOP, HSP40 and p23 (ref. 21). Given that Ago2 ligands are siRNA duplexes, RISC assembly is viewed analogously to steroid hormone receptor maturation. Thus, we asked whether RISC assembly could be reconstituted using fly Ago2, Dicer-2-R2D2 heterodimer and recombinant chaperone factors. As an *in vitro* assay system, we immunopurified Flag-Ago2 from *Drosophila* S2 cells (Extended Data Fig. 1a), and then incubated the immunoprecipitate with a siRNA duplex containing a ³²P-radiolabelled guide strand. On its own, immunopurified Ago2 failed to bind siRNA (Extended Data

Fig. 1b); when supplemented with lysate from *ago2*⁴¹⁴ mutant fly embryos (Extended Data Fig. 1c), Ago2 loaded siRNA duplex and ejected the passenger strand, forming mature Ago2-RISC containing single-stranded guide (Extended Data Fig. 1b). In contrast, mutant Ago2 bearing a Y897E substitution in the pocket that recognizes the phosphate and the base of the guide 5' nucleotide²² failed to load siRNA even in the presence of lysate (Extended Data Fig. 1b, d). A D965A catalytic mutant of Ago2 that cannot cleave the passenger strand²⁰ loaded siRNA duplex but failed to mature into Ago2-RISC (Extended Data Fig. 1b, d). As expected, duplex loading was blocked by chaperone inhibitors (Extended Data Fig. 1e and Supplementary Discussion).

To reconstitute RISC assembly, we first purified the three core components of the chaperone machinery—Hsc70-4, Hsp83 and Hop—using affinity-tagged Hop²³ (Extended Data Fig. 1a). Although Ago2-RISC assembly was undetectable with Dicer-2-R2D2 alone, supplementing Dicer-2-R2D2 with Hsc70-4, Hsp83 and Hop supported formation of pre- and mature Ago2-RISC (Fig. 1a and Extended Data Fig. 2a). Ago2-RISC assembly was further augmented by addition of recombinant Droj2 or p23 (Fig. 1a and Extended Data Figs 1a, 2a). To dissect the contribution of each of the seven proteins to the reaction, we individually purified recombinant Hsc70-4, Hsp83 and Hop (Extended Data Fig. 1a). Together, all seven proteins—Dicer-2-R2D2, Hsp83, Hsc70-4, Hop, Droj2 and p23—directed efficient Ago2-RISC assembly (Fig. 1b and Extended Data Fig. 2b). Omitting any of the five chaperone factors reduced the efficiency of RISC assembly by ~70–90% (Fig. 1b and Extended Data Fig. 2b). Reconstituted Ago2-RISC was as fully active for target cleavage as the Ago2-RISC programmed in lysate (Fig. 1c and Extended Data Fig. 2c–f). Thus, Dicer-2-R2D2 plus five chaperone proteins suffice to assemble siRNA duplexes into Ago2-RISC. In contrast to canonical duplex loading, loading of a single-stranded RNA into Ago2 bypassed the requirement of Dicer-2-R2D2 or the chaperones (Extended Data Fig. 2g, h).

Although it was reported that complementary target RNAs can promote release of guide strands²⁴, RISC is extremely stable in general. In fact, recombinant Ago proteins always co-purify with endogenous RNAs from the host cells^{25,26}, and our fly Ago2 preparation is largely preoccupied with endogenous siRNAs. If the chaperone machinery facilitates duplex loading, it may do so by releasing preloaded RNAs to increase the pool of empty Ago. However, loading of exogenous *let-7* siRNA duplex into fly Ago2 in the reconstituted system was not accompanied by the removal of the endogenous *esi-2.1* (also known as *esi-2*) RNA (Fig. 1d and Extended Data Fig. 3a) or the *in vitro* preloaded CXCR4 siRNA (Extended Data Fig. 3b–d). We conclude that the chaperones support an RNA-free form of Ago2 to load siRNA duplexes without emptying preassembled Ago2-RISC.

In the fly Ago2-RISC assembly pathway, Dicer-2-R2D2 senses the siRNA duplex thermodynamic asymmetry, thereby fixing the orientation of the duplex transfer to Ago2 (ref. 27). Accordingly, the strand with

¹Institute of Molecular and Cellular Biosciences, The University of Tokyo, Bunkyo-ku, Tokyo 113-0032, Japan. ²Department of Chemistry and Biotechnology, Graduate School of Engineering, The University of Tokyo, Bunkyo-ku, Tokyo 113-0033, Japan. ³Department of Medical Genome Sciences, Graduate School of Frontier Sciences, The University of Tokyo, Bunkyo-ku, Tokyo 113-0032, Japan. [†]Present address: Department of Molecular and Cell Biology, University of California, Berkeley, California, 94720, USA.

*These authors contributed equally to this work.

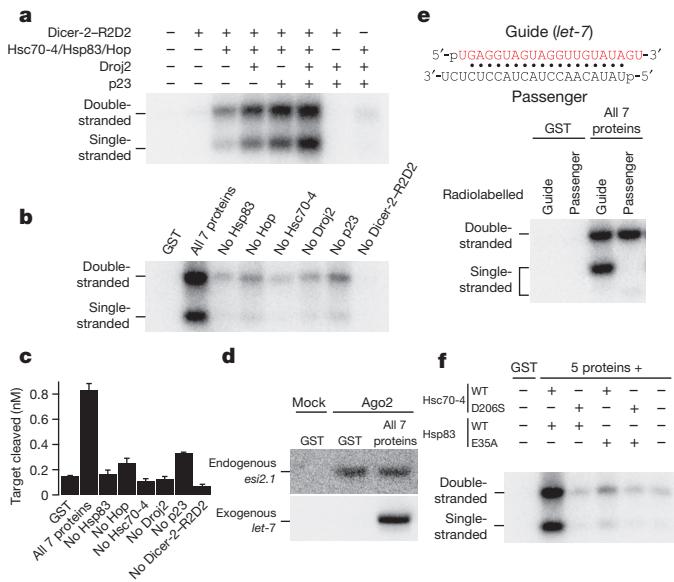


Figure 1 | In vitro reconstitution of Ago2–RISC assembly. **a**, Pull-down assay using the three core chaperone components and recombinant Dicer-2-R2D2, Droj2 and p23. **b**, Pull-down assay with the seven purified recombinant proteins. **c**, Target cleavage after 30-min reaction in the reconstituted system. Data represent mean \pm standard deviation (s.d.) ($n = 3$). **d**, Loading of exogenous *let-7* siRNA does not accompany unloading of endogenous *esi-2.1* RNA. RNAs in Ago2 were detected by northern blotting. **e**, The reconstituted system faithfully recapitulates the siRNA loading asymmetry. Ago2–RISC was assembled with an siRNA with either the guide or the passenger strand radiolabelled. **f**, ATP hydrolysis by both Hsc70-4 and Hsp83 is required for maximum Ago2–RISC assembly. WT, wild type.

the less thermodynamically stable 5' end is predominantly selected as the guide strand^{1,2}. When RISC assembly was initiated in the reconstituted system using a functionally asymmetric siRNA duplex with either the guide strand or the passenger strand radiolabelled, both strands were detected in pre-Ago2–RISC, but only the guide strand was detected in mature Ago2–RISC (Fig. 1e and Extended Data Fig. 4a). Thus, the reconstituted system faithfully recapitulates siRNA loading asymmetry.

ATP dependence is a hallmark of RISC assembly initiated by small RNA duplexes^{4–7,9,15}. In the seven-protein reconstituted system, the non-hydrolysable ATP analogues ATP- γ S, AMP-PNP and ADP all inhibited duplex loading (Extended Data Fig. 4b, c). Two of the five chaperone factors are ATPases (Supplementary Discussion), so we tested the ATPase-defective mutants Hsc70-4(D206S)²⁸ and Hsp83(E35A)²⁹. Both mutant proteins reduced duplex loading efficiency in the reconstituted system, suggesting that ATP hydrolysis by both Hsc70-4 and Hsp83 is required for maximum RISC assembly (Fig. 1f and Extended Data Fig. 4d).

Because the biochemical approach cannot detect intermediate states during the handover of the siRNA duplex from Dicer-2-R2D2 to Ago2, we performed single-molecule imaging of RISC assembly using total internal reflection fluorescence (TIRF) microscopy to dissect how the chaperone machinery assists duplex loading (Fig. 2a and Extended Data Fig. 5)³⁰. We expressed amino-terminal Halo-TEV-Flag-tagged fly Ago2, in which a cleavage site for tobacco etch virus (TEV) protease was inserted between the Halo and Flag tags, in S2 cells (Extended Data Fig. 6a). The Halo-TEV-Flag-tagged Ago2 was captured on a NeutrAvidin-derivatized quartz glass surface via Halo-ligand-conjugated biotin. After washing of the surface, we added the reconstituted system or equimolar glutathione S-transferase (GST) as a control, together with an siRNA with a 3' Alexa647 (red)-labelled guide strand and a 3' Alexa555 (green)-labelled passenger strand (Fig. 2a). The siRNA duplex with these 3'-fluor modifications efficiently formed pre- and mature Ago2–RISC *in vitro* (Extended Data Fig. 6b).

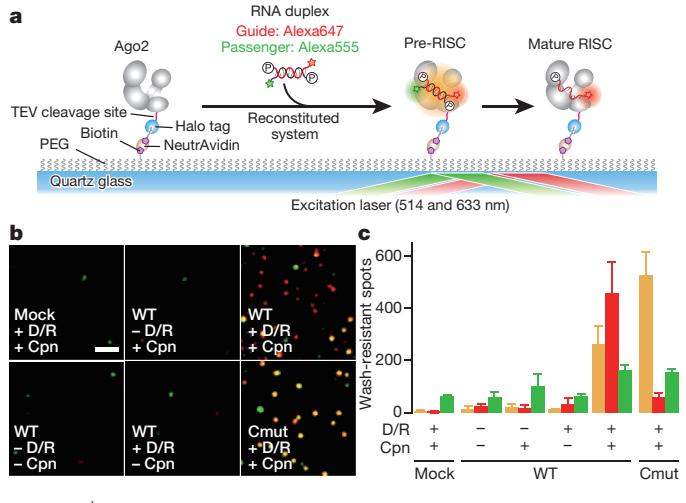


Figure 2 | Single-molecule observation of pre- and mature RISC formation. **a**, Schematic representation of the single-molecule analysis. **b, c**, Representative images (b) and quantification (c) of wash-resistant spots after 1 h incubation. Yellow, double-stranded signals; red, guide-alone signals; green, passenger-alone and background signals. Mock, no Ago2 tethered. Cpn, chaperone machinery; D/R, Dicer-2-R2D2; WT, wild type; Cmut, catalytic mutant (D965A). Scale bar, 2 μ m. Data represent mean \pm s.d. (WT, +D/R, +Cpn, $n = 6$; others, $n = 3$).

When the surface-tethered Ago2 and the seven-protein reconstituted system were present, many fluorescent spots were observed (Extended Data Fig. 6c); after 1 h incubation, the surface was washed and the remaining spots were analysed (Extended Data Fig. 5a and Supplementary Discussion). Two subpopulations were seen for guide-containing spots: one with the guide strand alone (red spots) and the other with both guide and passenger strands (colocalization of red and green signals, giving yellow spots) (Fig. 2b, c). These spots were observed only in the presence of Dicer-2-R2D2, the chaperone machinery and the surface-tethered Ago2; when Ago2 was released from the glass surface by TEV-protease-mediated cleavage (Extended Data Figs 5b and 6d) few yellow and red spots remained (Extended Data Fig. 6e, f). Moreover, the catalytic mutant Ago2, which is defective in strand separation (Extended Data Fig. 1b), predominantly produced yellow spots (Fig. 2b, c). These data suggest that red and yellow wash-resistant spots represent mature RISC and pre-RISC, respectively. In contrast, the number of green only spots was similar in any condition, suggesting that many of them are background signals that are commonly observed with green laser excitation (Fig. 2c). Supporting this idea, inverting the strand colours (that is, green-labelled guide and red-labelled passenger) resulted in accumulation of green spots in wild-type Ago2 and yellow spots in catalytic mutant Ago2 (Extended Data Fig. 6g, h) with few red spots. Our single-molecule assay, like the pull-down assay, thus faithfully recapitulates canonical RISC assembly.

Next, we continuously monitored the appearance of the spots during the first 20 min without washing (Extended Data Fig. 5c and Supplementary Discussion). Given that the siRNA duplex is loaded into Ago2 as double strands, we focused on the events where the red and green spots appeared at the same time. In many cases, the red and green spots disappeared simultaneously within the 20-min observation window (Fig. 3a, top; ~86%). In some cases, the two colours remained present until the end of the observation period (Fig. 3a, middle; ~9%), the green spots disappeared earlier than the red spots (Fig. 3a, bottom; ~3%), or vice versa (data not shown; ~1%). Interestingly, the number of total binding events was comparable in the presence or absence of the chaperone machinery, but required the presence of Dicer-2-R2D2 and surface-tethered Ago2 (Fig. 3b). Moreover, the distribution of the spot appearance duration showed a similar single-exponential function with a time constant (τ) of ~20 s, regardless of the presence of the chaperone machinery (Fig. 3c and Extended Data Fig. 7a). Thus, the chaperone

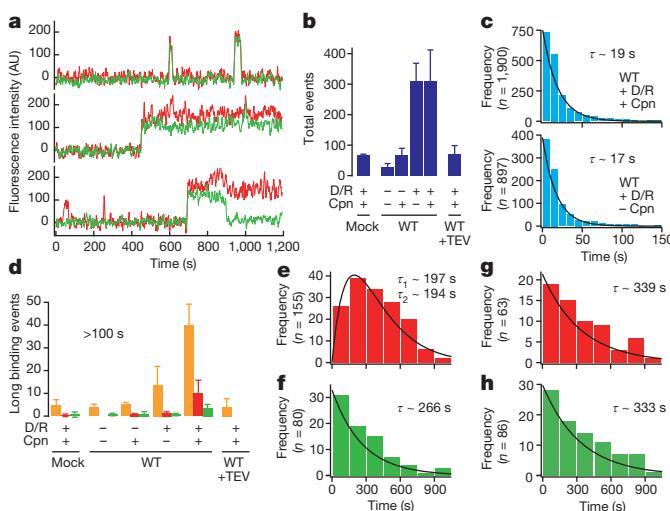


Figure 3 | The chaperone machinery extends the dwell time of Dicer-2-R2D2-siRNA on Ago2. **a**, Representative traces that start with simultaneous binding of red and green spots during the 20-min monitoring period. AU, arbitrary units. **b**, Total event numbers during the continuous monitoring. Dicer-2-R2D2 (D/R), but not the chaperone machinery (Cpn), was required for the spots to appear. Wild type (WT) + TEV, Ago2 was detached from the glass surface before monitoring (Extended Data Fig. 5c). **c**, Dwell time histograms. The chaperone machinery was dispensable for the transient binding ($\sim 20\text{ s}$). Solid line, fitted from the cumulative plot (Extended Data Fig. 7a). **d**, Numbers of long binding events ($>100\text{ s}$). Yellow, simultaneous dissociation of red and green spots or their continuous binding to the end of observation; red, dissociation of the green spots with the red spots remaining (yellow-to-red); green, dissociation of the red spots with the green spots remaining (yellow-to-green). WT + TEV, detachment of Ago2 before monitoring. **e–h**, Dwell-time distribution for the yellow-to-red or yellow-to-green transitions in wild-type (**e**, **f**) and catalytic mutant Ago2 (**g**, **h**). The yellow-to-red duration in wild-type Ago2 (**e**) fits to a two-state reaction model and differs significantly from those of the other data (**f–h**, $P = 0.00025, 0.024$, and 0.017 , respectively; two-sample Anderson–Darling test), which showed single-exponential decays. Solid line, fitted from the cumulative plot (Extended Data Fig. 7b–e). Data in **b**, **d** represent mean \pm s.d. (WT, +D/R, +Cpn, $n = 5$; others, $n = 3$).

machinery is dispensable for the short dwelling of Dicer-2-R2D2-bound siRNA on Ago2. In contrast, long binding events ($>100\text{ s}$) require Ago2, Dicer-2-R2D2 and the chaperone machinery (Fig. 3d). When all factors were present, $\sim 18\%$ of long binding events showed the disappearance of the green spots before the red spots, which presumably reflects the formation of mature RISC via ejection of the passenger strand from Ago2 (Fig. 3d, red).

To determine the time constant for RISC maturation, we further analysed the spots that displayed the yellow-to-red transition. The distribution of the dwell time before the disappearance of the green spots showed a clear peak and fitted well to a two-step model with time constants of $\sim 197\text{ s}$ and $\sim 194\text{ s}$ (Fig. 3e and Extended Data Fig. 7b), indicating the involvement of two or more consecutive first-order reactions. This is in line with the model that passenger-strand ejection occurs only after the initial, transient binding between Dicer-2-R2D2-siRNA and Ago2, extension of the dwell time by the chaperone machinery, and the cleavage of the passenger strand. In contrast, the spots that displayed the opposite, yellow-to-green transition (Fig. 3d, green), showed a single-exponential decay of the duration (Fig. 3f and Extended Data Fig. 7c), presumably reflecting the photobleaching of the red fluorophore, rather than ‘flipped’ RISC maturation. Moreover, in the catalytic mutant Ago2, which should retain the guide and passenger strands, both red and green spots showed single-exponential decays (Fig. 3g, h and Extended Data Fig. 7d, e). We conclude that the chaperone machinery acts to extend the dwell time of the Dicer-2-R2D2-siRNA complex on Ago2, thereby facilitating the siRNA handover from Dicer-2-R2D2 to Ago2 that leads to RISC maturation.

The 5'-phosphate on the guide strand is thought to be crucial for Ago2 to anchor the guide strand via the 5' nucleotide-binding pocket^{25,26}. Indeed, a 5'-hydroxyl group on the guide strand, which did not affect siRNA binding to Dicer-2-R2D2 (Supplementary Discussion), blocked formation of pre- and mature Ago2–RISC without altering the functional asymmetry of the two strands (Extended Data Fig. 8a–c). This was also confirmed by the single-molecule analysis; we observed few wash-resistant spots when the fluorescently labelled siRNA duplex bore a 5'-hydroxyl group on the guide strand or when the 5' nucleotide-binding pocket of Ago2 was mutated (Fig. 4a, b). Intriguingly, however, continuous monitoring without washing showed that the number of total events for the 5'-hydroxyl siRNA or the 5' nucleotide-binding pocket mutant of Ago2 was comparable with that for the 5'-phosphate siRNA and wild-type Ago2 (Fig. 4c). The distribution of the dwell time fitted to similar single exponential curves with a time constant of $\sim 20\text{ s}$, regardless of recognition of the 5'-phosphate (Figs 3c, 4d and Extended Data Figs 7a, 8d). In contrast, the long binding events that lead to RISC maturation were much less frequent in the absence of 5'-phosphate recognition (Fig. 4e). Thus, 5'-phosphate recognition by Ago2 is dispensable for the transient binding between Dicer-2-R2D2-siRNA and Ago2, but is required for the extension of the dwell time.

Our data define the following multiple fundamental steps in Ago2–RISC assembly (Extended Data Fig. 9). A Dicer-2-R2D2-bound siRNA duplex associates with but rapidly dissociates from unchaperoned Ago2. The chaperone machinery extends the dwell time of the Dicer-2-R2D2-siRNA complex on Ago2 to facilitate siRNA handover, but this extension requires the recognition of the guide 5'-phosphate by Ago2 (Extended Data Fig. 9). Previous studies suggest that, although each individual domain of Ago appears structurally rigid, the overall domain organization of an RNA-free state of Ago is highly flexible, to which small RNA binding confers remarkable stability^{16,20,25,26}. We envision that the Hsc70/Hsp90 chaperone machinery has a critical role in supporting a loading-competent state of Ago2, in which the 5' nucleotide-binding

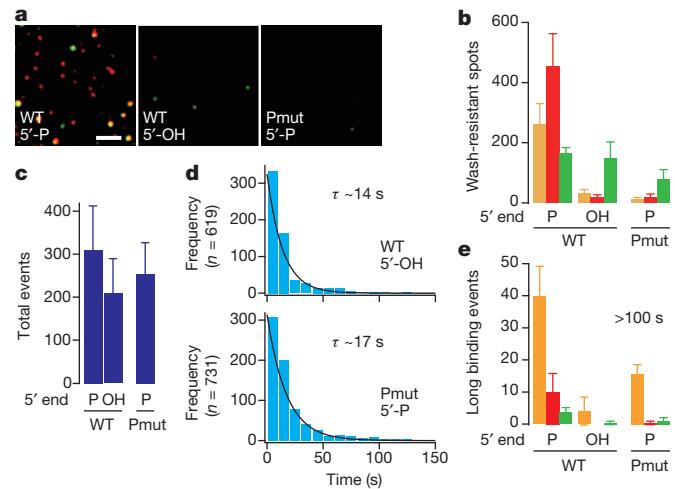


Figure 4 | Recognition of the guide 5'-phosphate is required for dwell-time extension. **a, b**, Representative images (**a**) and quantification (**b**) of wash-resistant spots after incubation. 5'-Phosphate recognition was required for the formation of pre- and mature RISC. 5'-OH, 5'-hydroxyl group on the guide strand; 5'-P, 5'-phosphate; Pmut, the 5' nucleotide-binding pocket mutant (Y897E); WT, wild type. Scale bar, $2\mu\text{m}$. The wild-type:5'-phosphate data are the same as in Fig. 2b, c. **c**, Total event numbers during continuous monitoring. 5'-Phosphate recognition was not necessary for the spots to appear. The wild-type:5'-phosphate data are the same as in Fig. 3b. **d**, Dwell time histograms. 5'-Phosphate recognition was not necessary for transient binding. Solid line, fitted from the cumulative plot (Extended Data Fig. 8d). **e**, Numbers of long binding events ($>100\text{ s}$). 5'-Phosphate recognition was required for dwell-time extension. The wild-type:5'-phosphate data are the same as in Fig. 3d. Data in **b**, **c**, **e** represent mean \pm s.d. ($n = 3$).

pocket is accessible to the 5'-phosphate of the guide strand of the Dicer-2-R2D2-bound siRNA duplex (Extended Data Fig. 9). Proper anchoring of the 5'-phosphate would trigger loading of the rest of the duplex, leading to successful RISC maturation coupled with closing and stabilization of Ago2. A future challenge will be to monitor directly a series of such dynamic conformational changes of Ago proteins.

Online Content Methods, along with any additional Extended Data display items and Source Data, are available in the online version of the paper; references unique to these sections appear only in the online paper.

Received 4 July 2013; accepted 23 January 2015.

Published online 30 March; corrected online 27 May 2015 (see full-text

HTML version for details).

1. Carthew, R. W. & Sontheimer, E. J. Origins and mechanisms of miRNAs and siRNAs. *Cell* **136**, 642–655 (2009).
2. Kawamata, T. & Tomari, Y. Making RISC. *Trends Biochem. Sci.* **35**, 368–376 (2010).
3. Liu, Q. *et al.* R2D2, a bridge between the initiation and effector steps of the *Drosophila* RNAi pathway. *Science* **301**, 1921–1925 (2003).
4. Pham, J. W., Pellino, J. L., Lee, Y. S., Carthew, R. W. & Sontheimer, E. J. A. Dicer-2-dependent 80S complex cleaves targeted mRNAs during RNAi in *Drosophila*. *Cell* **117**, 83–94 (2004).
5. Tomari, Y. *et al.* RISC assembly defects in the *Drosophila* RNAi mutant armitage. *Cell* **116**, 831–841 (2004).
6. Iki, T. *et al.* In vitro assembly of plant RNA-induced silencing complexes facilitated by molecular chaperone HSP90. *Mol. Cell* **39**, 282–291 (2010).
7. Iwasaki, S. *et al.* Hsc70/Hsp90 chaperone machinery mediates ATP-dependent RISC loading of small RNA duplexes. *Mol. Cell* **39**, 292–299 (2010).
8. Miyoshi, T., Takeuchi, A., Siomi, H. & Siomi, M. C. A direct role for Hsp90 in pre-RISC formation in *Drosophila*. *Nature Struct. Mol. Biol.* **17**, 1024–1026 (2010).
9. Kawamata, T., Seitz, H. & Tomari, Y. Structural determinants of miRNAs for RISC loading and slicer-independent unwinding. *Nature Struct. Mol. Biol.* **16**, 953–960 (2009).
10. Kanelloupolou, C. *et al.* Dicer-deficient mouse embryonic stem cells are defective in differentiation and centromeric silencing. *Genes Dev.* **19**, 489–501 (2005).
11. Murchison, E. P., Partridge, J. F., Tam, O. H., Cheloufi, S. & Hannon, G. J. Characterization of Dicer-deficient murine embryonic stem cells. *Proc. Natl. Acad. Sci. USA* **102**, 12135–12140 (2005).
12. Ye, X. *et al.* Structure of C3PO and mechanism of human RISC activation. *Nature Struct. Mol. Biol.* **18**, 650–657 (2011).
13. Betancur, J. G. & Tomari, Y. Dicer is dispensable for asymmetric RISC loading in mammals. *RNA* **18**, 24–30 (2012).
14. Kim, Y. *et al.* Deletion of human *tarbp2* reveals cellular microRNA targets and cell-cycle function of TRBP. *Cell Rep.* **9**, 1061–1074 (2014).
15. Yoda, M. *et al.* ATP-dependent human RISC assembly pathways. *Nature Struct. Mol. Biol.* **17**, 17–23 (2010).
16. Johnston, M., Geoffroy, M. C., Sobala, A., Hay, R. & Hutvagner, G. HSP90 protein stabilizes unloaded argonaute complexes and microscopic P-bodies in human cells. *Mol. Biol. Cell* **21**, 1462–1469 (2010).
17. Matranga, C., Tomari, Y., Shin, C., Bartel, D. P. & Zamore, P. D. Passenger-strand cleavage facilitates assembly of siRNA into Ago2-containing RNAi enzyme complexes. *Cell* **123**, 607–620 (2005).
18. Miyoshi, K., Tsukumo, H., Nagami, T., Siomi, H. & Siomi, M. C. Slicer function of *Drosophila* Argonautes and its involvement in RISC formation. *Genes Dev.* **19**, 2837–2848 (2005).
19. Rand, T. A., Petersen, S., Du, F. & Wang, X. Argonaute2 cleaves the anti-guide strand of siRNA during RISC activation. *Cell* **123**, 621–629 (2005).
20. Liu, Y. *et al.* C3PO, an endoribonuclease that promotes RNAi by facilitating RISC activation. *Science* **325**, 750–753 (2009).
21. Smith, D. F. & Toft, D. O. The intersection of steroid receptors with molecular chaperones: observations and questions. *Mol. Endocrinol.* **22**, 2229–2240 (2008).
22. Rüdel, S. *et al.* Phosphorylation of human Argonaute proteins affects small RNA binding. *Nucleic Acids Res.* **39**, 2330–2343 (2011).
23. Dittmar, K. D. & Pratt, W. B. Folding of the glucocorticoid receptor by the reconstituted Hsp90-based chaperone machinery. The initial hsp90/p60/hsp70-dependent step is sufficient for creating the steroid binding conformation. *J. Biol. Chem.* **272**, 13047–13054 (1997).
24. De, N. *et al.* Highly complementary target RNAs promote release of guide RNAs from human Argonaute2. *Mol. Cell* **50**, 344–355 (2013).
25. Nakanishi, K., Weinberg, D. E., Bartel, D. P. & Patel, D. J. Structure of yeast Argonaute with guide RNA. *Nature* **486**, 368–374 (2012).
26. Schirle, N. T. & MacRae, I. J. The crystal structure of human Argonaute2. *Science* **336**, 1037–1040 (2012).
27. Tomari, Y., Matranga, C., Haley, B., Martinez, N. & Zamore, P. D. A protein sensor for siRNA asymmetry. *Science* **306**, 1377–1380 (2004).
28. Flaherty, K. M., DeLuca-Flaherty, C. & McKay, D. B. Three-dimensional structure of the ATPase fragment of a 70K heat-shock cognate protein. *Nature* **346**, 623–628 (1990).
29. Prodromou, C. *et al.* Identification and structural characterization of the ATP/ADP-binding site in the Hsp90 molecular chaperone. *Cell* **90**, 65–75 (1997).
30. Miyazono, Y., Hayashi, M., Karagiannis, P., Harada, Y. & Tadakuma, H. Strain through the neck linker ensures processive runs: a DNA-kinesin hybrid nanomachine study. *EMBO J.* **29**, 93–106 (2010).

Supplementary Information is available in the online version of the paper.

Acknowledgements We thank Q. Liu for providing vectors for Dicer-2 and R2D2 expression, M. Horwich and P. D. Zamore for Ago2 cDNA, and H. Siomi and M. C. Siomi for anti-Ago2 antibody and ago2⁴¹⁴ flies. We are also grateful to P. B. Kwak and A. Tsutsumi for assistance with plasmid constructions, S. Katsuma for support for Dicer-2-R2D2 preparation, and A. Yamashita for advice on protein purification. We thank H. Taguchi and T. Ueda for fruitful discussions and P. D. Zamore, H. Seitz and the members of the Tomari laboratory for critical comments on the manuscript. This work was supported in part by Grants-in-Aid for Scientific Research on Innovative Areas ('Functional machinery for non-coding RNAs' and 'Non-coding RNA neo-taxonomy'), a Grant-in-Aid for Young Scientists (A) (to H.T. and Y.T.), a Grant-in-Aid for Research Activity start-up (to S.I.), and Grants-in-Aid for challenging Exploratory Research (to H.M.S.) from The Ministry of Education, Culture, Sports, Science and Technology in Japan.

Author Contributions S.I. performed biochemical experiments and H.M.S. performed single-molecule experiments. Y.S. and T.S. performed mass spectrometry analysis. S.I., H.M.S., H.T. and Y.T. designed experiments, analysed data and wrote the manuscript.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to Y.T. (tomari@iam.u-tokyo.ac.jp) or H.T. (tadakuma@k.u-tokyo.ac.jp).

METHODS

General methods. Lysis buffer, lysates from *Drosophila melanogaster* embryos and S2 cells, radiolabelled small RNA duplexes, formamide dye, and cap-labelled target RNAs were prepared as previously described^{31–35}. Native gel analysis of double-stranded and single-stranded RNAs and gel shift assay were performed as previously described^{31–35}. Small RNAs shown in Fig. 1e and Extended Data Fig. 4a were purchased from GeneDesign. For single-molecule analysis, we used 5'-pUGAGG UAGUAGGUUGUAUAGU-3' (*let-7*, guide strand), 5'-UGAGGUAGUAGGU UGUUAUAGU-3' (*let-7*, guide strand, 5'-OH), and 5'-pUAUACAACCUACUAC CUCUCU-3' (passenger strand) with chemical labelling of Alexa647 or 555 through a C6-amino linker at the 3' end purchased from GeneDesign. The siRNA duplexes bore a single unpaired nucleotide at the 5' end of the guide strand, making it highly asymmetric.

Plasmid constructions. *pCold-Hsp83*, *Hsp83(E35A)*, *Hsc70-4*, *Hsc70-4(D206S)*, *Hop*, *Droj2* and *p23*. DNA fragments containing *hsp83*, *hop*, *droj2*, *p23* (*cg16817*), *hsc70-4* or *hsc70-4(D206S)* genes were amplified by PCR from S2 cell cDNA, pENTR WT *Hsc70-4* or pENTR *Hsc70-4(D206S)* (ref. 7) and cloned into *pColdII* (Takara) with NdeI-EcoRI or NdeI-HindIII restriction enzyme sites. E35A mutation in *hsp83* gene was introduced by QuickChange site-directed mutagenesis (Stratagene).

pASW-Hop. A DNA fragment containing the *hop* gene was amplified from S2 cell cDNA and cloned into pENTR/D-TOPO (Invitrogen) according to the manufacturer's instruction, followed by recombination with *pASW* destination vector⁷ with Gateway LR Clonase II (Invitrogen).

pAFW-Ago2 wild type, *D965A* and *Y897E*. *pAFW-Ago2* wild type was described previously⁷. *D965A* and *Y897E* mutations were introduced by QuikChange site-directed mutagenesis (Stratagene).

pAHaloFW-Ago2 wild type, *D965A* and *Y897E*. A DNA fragment containing the Halo tag and TEV protease recognition sequence from pFN18A (Promega) was inserted into *pAFW* (*Drosophila* Gateway vector collection, Actin 5C promoter and N-terminal 3×Flag tag) at the 5' end upstream of the 3×Flag-tag sequence by In-Fusion HD (Clontech), to construct the *pAHaloFW* destination vector. *pENTR-Ago2* wild type was a gift from M. Horwich and P. D. Zamore. *pDONR-Ago2 D965A* and *Y897E* were constructed by recombination with *pDONR/Zeo* (Invitrogen) and *pAFW-Ago2 D965A* or *Y897E* by BP Clonase II (Invitrogen). *pENTR-Ago2* wild type, *pDONR-Ago2 D965A* and *Y897E* were recombined with *pAHaloFW* to construct *pAHaloFW-Ago2* wild type, *D965A* and *Y897E*, respectively.

pAWH-Dicer-2 and *pASW-R2D2*. DNA fragments containing the *dicer-2* or *r2d2* gene were amplified by PCR from *pFastBac-Dicer-2* or *pFastBac-R2D2* (refs 3, 27), which were gifts from Q. Liu, and cloned into pENTR/D-TOPO (Invitrogen) or pENTR/TEV/D-TOPO (Invitrogen), followed by recombination with *pAWH* (*Drosophila* Gateway vector collection) or *pASW* destination vector⁷ with Gateway LR Clonase II (Invitrogen).

DNA oligonucleotides for plasmid construction are as follows. *Hsp83*, 5'-AAA CCCCATATGATGCCAGAAGCAGAGA-3' and 5'-AAACCCGAATTCTT ATTCGACCTCCTCCATG-3'; *Hsc70-4*, 5'-AAACCCCATATGATGTCTAAA GCTCCTGCTG-3' and 5'-AAACCCGAATTCTTAGTCGACCTCTCGATG-3'; *Hop* for *pColdI*, 5'-AAACCCCATATGATGGACAAGGTGAACGAAC-3' and 5'-AAACCCGAATTCTTAGTCGAGATCTGAATGATGCC-3'; *Droj2*, 5'-AAACC CCATATGATGGTTAAGGAGACTGGATA-3' and 5'-AAACCCGAATTCTTA ACTCGATGTGCACTGC-3'; *p23*, 5'-AAACCCCATATGATGTCGGCA GCAGCAGG-3' and 5'-AAACCCAAGCTCTAGGCAGCTGGCTCTT-3'; *Hsp83 E35A*, 5'-GGAGATTTCTCGCGCGTTGATCTCGAACGCTTC-3' and 5'-G AAGCGTTGAGATCAACGCGCGCAGGAAAATCTCC-3'; *Hop* for *pENTR*, 5'-CACCATGGACAAGGTGAACGAAC-3' and 5'-CTAGTGAACTGATG ATGCC-3'; *Ago2 D965A*, 5'-ACACCATGTCATTGGAGGCCGCTGACCC ATCCCTCTCCCGA-3' and 5'-TCGGGAGAGGGATGGGTCA CGCGCTCAATGTACATGGTGT-3'; *Ago2 Y897E*, 5'-ATT CCTCAGTTAGAATTTCT GAAGATACAATTAAGCAGAAGGCC-3' and 5'-GGCCTCTGCTTAATTG TATCTTCAGAAATTCTAAACTGAGGAAT-3'; *Dicer-2*, 5'-CACCATGGAA GATGTGGAAATCAA-3' and 5'-TTAGGCGTCGCATTTGCT-3'; and *R2D2*, 5'-CACCATGGATAACAAAGTCAGCCG-3' and 5'-TTAAATCAACATGGTG CGAAA-3'.

Protein purifications. Recombinant proteins including *Hsp83*, *Hsp83(E35A)*, *Hop*, *Hsc70-4*, *Hsc70-4(D206S)*, *Droj2* and *p23* were expressed as N-terminal His-tagged proteins in *Escherichia coli* BL21 strain. Typically, the cells in 2-litre cultures were cultivated to an OD_{600 nm} of 0.5 at 37 °C with 100 µg ml⁻¹ ampicillin and then grown at 15 °C overnight with 1 mM isopropyl β-D-thiogalactoside (IPTG). The cell pellets were resuspended in His A buffer (30 mM HEPES-KOH pH 7.4, 200 mM KOAc, 2 mM Mg(OAc)₂, 5% glycerol, 20 mM imidazole, 0.2 mM TCEP) containing 1× EDTA-free protease inhibitor cocktail (Roche), sonicated, and centrifuged at 10,000g for 20 min. The supernatant was loaded onto HisTrap FF crude

5 ml (GE Healthcare) and eluted with linear gradient from His A buffer to His B buffer (His A buffer containing 400 mM imidazole). The peak fractions were collected, diluted ten times with 30 mM HEPES-KOH pH 7.4, 2 mM Mg(OAc)₂, 5% glycerol, 1 mM dithiothreitol (DTT), loaded onto MonoQ 1 ml (GE Healthcare), and eluted with linear gradient from MonoQ A buffer (30 mM HEPES-KOH pH 7.4, 20 mM KCl, 2 mM Mg(OAc)₂, 5% glycerol, 1 mM DTT) to MonoQ B buffer (MonoQ A buffer containing 1 M KCl). The peaked fractions were collected, buffer-exchanged to 30 mM HEPES-KOH pH 7.4, 100 mM KOAc, 2 mM Mg(OAc)₂, 10% glycerol, 1 mM DTT with PD-10 (GE Healthcare). For *Hsp83* recombinant protein purification, glycerol and reducing reagent were omitted from all the procedures. The expression and purification of the complex of streptavidin-binding peptide (SBP)-Hop, endogenous *Hsp83*, and endogenous *Hsc70-4* were performed according to the method previously described with some modifications^{36,37}. Briefly, S2 cells (1.0 × 10⁶ cells ml⁻¹) were transfected with 10 µg *pASW-Hop* and 20 µl X-tremeGENE HP (Roche) per 10 ml and cultured for 72 h. Transfected S2 cells were collected and washed once with PBS. Cell pellets were resuspended with twice the volume of hypotonic lysis buffer (10 mM HEPES-KOH pH 7.4, 10 mM KCl, 1.5 mM MgCl₂, 1× EDTA-free protease inhibitor cocktail (Roche)), incubated for 30 min, and centrifuged at 17,000g for 20 min. Typically, 2 ml supernatant was incubated with 200 µl Streptavidin Sepharose HP (GE Healthcare) for 1 h. The beads were washed five times with lysis buffer and SBP-Hop and its binding proteins were eluted with 100 µl lysis buffer containing 5 mM biotin for 30 min. The proteins around 85 kDa and 70 kDa (Extended Data Fig. 1a) were identified as *Hsp83* and *Hsc70-4*, respectively, by mass spectrometry as previously described⁷. Recombinant Dicer-2-R2D2 heterodimer was prepared by plasmid-directed expression of SBP-R2D2 and Dicer-2 in S2 cells^{36,37}. 7.5 µg *pAWH-Dicer-2* and 7.5 µg *pASW-R2D2* were co-transfected per 10 ml of S2 cell culture by 30 µl X-tremeGENE HP (Roche). After incubation with lysate, the beads were washed five times with lysis buffer containing 0.8 M NaCl, and rinsed twice with lysis buffer containing 30% glycerol and 1 mM DTT. The Dicer-2-SBP-R2D2 heterodimer was eluted with lysis buffer containing 2.5 mM biotin, 30% glycerol and 1 mM DTT. All purification steps were performed at 4 °C. Column chromatography was performed using AKTA purifier (GE Healthcare). All purified proteins were shock-frozen in liquid nitrogen and stored at -80 °C.

Small RNA pull-down, target cleavage and northern blot analysis. Typically, for 15 reactions, 30 µl Dynabeads Protein G (Invitrogen) were equilibrated with lysis buffer, incubated with 1 µl of 1 mg ml⁻¹ anti-Flag M2 antibody (F1804, Sigma) at 4 °C for 30 min, washed three times with lysis buffer, and incubated at 4 °C for 1 h with 30 µl lysate from S2 cells in which *pAFW-Ago2* wild type, *D965A* or *Y897E* was expressed. For Extended Data Fig. 3b–d, 5 µl lysate was pre-incubated with 5 nM ³²P-radiolabelled CXCR4 siRNAs, 1 mM ATP, 25 mM creatine monophosphate (Fluka), 0.03 U µl⁻¹ creatine kinase (Cal-biochem), 0.1 U µl⁻¹ RNasin Plus (Promega) and lysis buffer in 10 µl at 25 °C for 2 h. The beads were washed five times with wash buffer (lysis buffer containing 0.8 M NaCl and 1% Triton X-100) and rinsed with lysis buffer. Typically, a reaction contained approximately 40 fmol Ago2 protein. The beads were incubated in lysate or the reconstituted system with 20 nM ³²P-radiolabelled small RNA duplex, 1 mM ATP, 25 mM creatine monophosphate, 0.03 U µl⁻¹ creatine kinase, 0.1 U µl⁻¹ RNasin Plus and lysis buffer in 10 µl at 25 °C for 90 min. Five microlitres (Extended Data Fig. 1b, e) and 1 µl (Extended Data Fig. 2d–f) of *ago2*¹¹⁴ embryo lysate were used. PES (Sigma) and Celastrol (Sigma) were dissolved in dimethylsulphoxide (DMSO) and added into the reactions in Extended Data Fig. 1e, providing the final concentration of 3 mM PES and 3 mM Celastrol with 3% DMSO. Twenty nanomolar Dicer-2-R2D2, 2.2 µM *Hsp83-Hop-Hsc70-4* complex, 8.4 µM *Droj2* and 2.2 µM *p23* were used in Fig. 1a, d and Extended Data Figs 3, 6b and 8b, c. Twenty nanomolar Dicer-2-R2D2, 4.4 µM *Hsp83* or *Hsp83(E35A)*, 2.2 µM *Hsc70-4* or *Hsc70-4(D206S)*, 2.2 µM *Hop*, 8.4 µM *Droj2* and 2.2 µM *p23* were used in Fig. 1b, c, e, f and Extended Data Fig. 2. In Fig. 1a–c, f and Extended Data Figs 2a–c, g, h and 4d, an equimolar GST to the omitted protein was supplemented to normalize the total protein concentration in the reactions. In Extended Data Fig. 2g, h, 20 nM ³²P-radiolabelled single-stranded guide strand was used with 100 µg ml⁻¹ yeast tRNA (Sigma). After incubation, the beads were washed four times with wash buffer. Bound RNAs were purified by Proteinase K treatment and run on 12% native polyacrylamide gel to separate double-stranded and single-stranded RNAs as previously described³¹. The protein levels of Flag-Ago2, Dicer-2, R2D2 and the chaperone components were unaffected by PES, Celastrol, non-hydrolysable ATP analogues or omission or mutation of a component(s). For cleavage assay, beads were further rinsed with lysis buffer twice and incubated with 5 nM cap-labelled 182-nucleotide target RNA for Fig. 1c and Extended Data Fig. 2c, or 5 nM cordycepin-labelled 28-nucleotide target RNA (5'-GAUACUAUACAACCUACUACCUACCU-3') for Extended Data Fig. 2e in lysis buffer containing 4 mM Mg(OAc)₂. The reaction was quenched by formamide dye and run on 8% and 15% denaturing polyacrylamide gel for 182-nucleotide and 28-nucleotide target RNA, respectively. Estimation of RISC concentration by pre-

steady-state 'burst' kinetics was performed as previously described³⁸. For northern blot, RNAs were run on 15% denaturing gel, transferred and chemically cross-linked to Hybond-N (GE Healthcare) by 1-ethyl-3-(3-dimethylaminopropyl) carbodiimide (EDC)^{39,40}. RNA probe 5'-GGAGCGAACUUGUUGGAGUCAA-3' for *esi2.1* and 2'-O-Me probe 5'-UCUUCACAUACAAACCUACUACCUAACC UU-3' for *let-7* were labelled by PNK (Takara) at the 5' end and hybridized with PerfectHyb plus (Sigma) at 45 °C and 68 °C, respectively. The images were acquired by a FLA-7000 imaging system (Fujifilm Life Sciences).

Western blot analysis. Anti-Flag M2 (F1804, Sigma) (1:5,000), anti-tubulin (B-5-1-2, Sigma) (1:10,000), anti-Ago2 (ref. 18) (1:100) and anti-HaloTag (G9211, Promega) (1:1,000) antibodies were used as the primary antibody. Signal Enhancer HIKARI (Nacalai tesque) was used for Ago2 detection. Chemiluminescence was induced by SuperSignal West Dura substrate (Pierce) or Luminata Forte Western HRP Substrate (Millipore) and images were acquired by a LAS-3000 imaging system (Fujifilm Life Sciences).

Protein labelling with Halo-tag ligands. For visualization of Halo-tagged protein (Extended Data Fig. 6a), 2 µl lysate from S2 cells, in which pAHaloFW-Ago2 wild type, D965A or Y897E was expressed, was incubated with 1 µM HaloTag TMR ligand (G8251, Promega) and 10% DMSO in 10 µl of lysis buffer at 25 °C for 30 min. After SDS-PAGE, the labelled proteins were visualized by a LAS-3000 image system. For single-molecule analysis, 400 µl fresh lysate was incubated with 1 µM HaloTag PEG-biotin ligand (G8591, Promega) and 1% DMSO at 25 °C for 30 min. Then, free ligands were removed by NAP-5 column (GE Healthcare) and equilibrated with lysis buffer containing 1 mM DTT. Flow-through fraction (~480 µl) was divided into tubes, 10 µl each, shock-frozen in liquid nitrogen, and stored at -80 °C.

Preparation of the observation chamber for single-molecule imaging. The observation chamber was prepared essentially as previously described^{30,41}. Halo-Ago2 was immobilized on the surface as illustrated in Fig. 2a and Extended Data Fig. 5. Typically, 20 µl of diluted S2 cell lysate containing biotin-labelled Halo-Ago2 or naïve S2 cell lysate as a mock was infused into the flow chamber. Then, the glass surface was washed with 50 µl of wash buffer twice, rinsed with 50 µl of lysis buffer three times, and 20 µl of lysis buffer containing 1% Biolipidure-203 (NOF Corporation). All experiments were performed at 25 ± 2 °C.

Reaction mixtures and buffers for single-molecule imaging. For snapshot observation, the reaction mixture contained 20 nM siRNA duplex, of which each strand was labelled with 3' Alexa fluorescent dye, 20 nM Dicer-2-R2D2, 1.2 µM Hsp83-Hop-Hsc70-4 complex, 0.75 µM Drok2, 2.7 µM p23, 1 mM ATP, 25 mM creatine monophosphate, 0.03 U µl⁻¹ creatine kinase, 0.1 U µl⁻¹ RNasin Plus, O₂ scavenger system (6.25 U ml⁻¹ glucose oxidase (G2133, Sigma), 62.5 U ml⁻¹ catalase (C40, Sigma) and 5.63 mg ml⁻¹ glucose), 1% Biolipidure-203 and 2 mM Trolox (238813, Sigma) in 20 µl of lysis buffer. For continuous monitoring, the reaction mixture contained 10 nM siRNA duplex, of which each strand was labelled with 3' Alexa fluorescent dye, 20 nM Dicer-2-R2D2, 1.2 µM Hsp83-Hop-Hsc70-4 complex, 0.75 µM Drok2, 2.7 µM p23, 1 mM ATP, 25 mM creatine monophosphate, 0.03 U µl⁻¹ creatine kinase, 0.1 U µl⁻¹ RNasin Plus, O₂ scavenger system, 1% Biolipidure-203, 50 µg ml⁻¹ yeast tRNA (R5636, Sigma) and 2 mM Trolox in 20 µl of lysis buffer. The observation mixture contained O₂ scavenger system and 2 mM Trolox in 20 µl of lysis buffer.

General image acquisition procedures. Single-molecule images were visualized by a total internal reflection fluorescence microscope equipped on an inverted type microscope (IX71, Olympus) with a ×150 oil immersion objective lens (UAPON 150 × OTIRFM, NA 1.45, Olympus), as previously described^{30,41}. The Alexa555 and 647 dyes were illuminated simultaneously with an argon (Ar) laser (514 nm; 35LAP321, Melles Griot) and a helium-neon (He-Ne) laser (633 nm; GLG5410, SOC), respectively. Fluorescence images from Alexa555 and 647 were separated by using a DualView2 (Optical Insights) and then projected side-by-side onto a back-illuminated electron-multiplying charge-coupled device (EMCCD) camera (iXon3 DU-897E-CSO-#BV, 512 × 512 pixels, Andor Technology). Images are pseudo-coloured in green (Alexa555) or red (Alexa647) in the figures.

General data analysis procedures for single-molecule imaging. Images were analysed using Image J software (<http://rsb.info.nih.gov/ij/>) with a built-in function and custom-designed plug-in software. The fluorescent intensity of the spots was measured using 6-pixels diameter circular regions of interest (ROIs), and only spots with fluorescence intensities greater than 3σ of the background intensity were analysed in order to exclude shot noise. All graphs and fitting curves were generated using KaleidaGraph (Synergy Software) and Igor Pro (WaveMetrics).

Measuring the steps and time constants of photobleaching. For measuring the steps and time constants of photobleaching (Extended Data Fig. 10a-c), 10 µl lysate from S2 cells, in which pAHaloFW-Ago2 D965A was expressed, was pre-incubated with 20 nM siRNA duplex, of which the guide strand was labelled with 3' Alexa647 and the passenger strand with 3' Alexa555, 1 mM ATP, 25 mM creatine monophosphate, 0.03 U µl⁻¹ creatine kinase, 0.1 U µl⁻¹ RNasin Plus and lysis buffer in 20 µl

at 25 °C for 30 min. The mixture was then incubated with 1 µM HaloTag PEG-biotin ligand and 1% DMSO at 25 °C for 30 min, and diluted with lysis buffer to 50 µl. Free ligands were removed by G-25 spin column (Invitrogen) equilibrated with lysis buffer. For measuring the steps of photobleaching, the flow-through fraction was diluted ~10-fold before infusion. The flow-through fraction was infused into the observation chamber, and after washing the chamber by wash buffer, the observation mixture was infused into the chamber. Then six random and non-overlapping places were chosen and imaged continuously for 600–1,800 s at frame rate of 1 frames s⁻¹ with either of the excitation lasers at three power densities (160, 320 and 640 mW mm⁻² for the Ar laser and 3.125, 6.25 and 12.5 W mm⁻² for the He-Ne laser). Three independent chambers were used for imaging. For measuring time constants of photobleaching, the mean fluorescence in each image was plotted versus time, and the resulting curve yielded a time constant of the photobleaching of each condition. We then plot the photobleaching time constants for Alexa555 and 647 versus the inverse of power density.

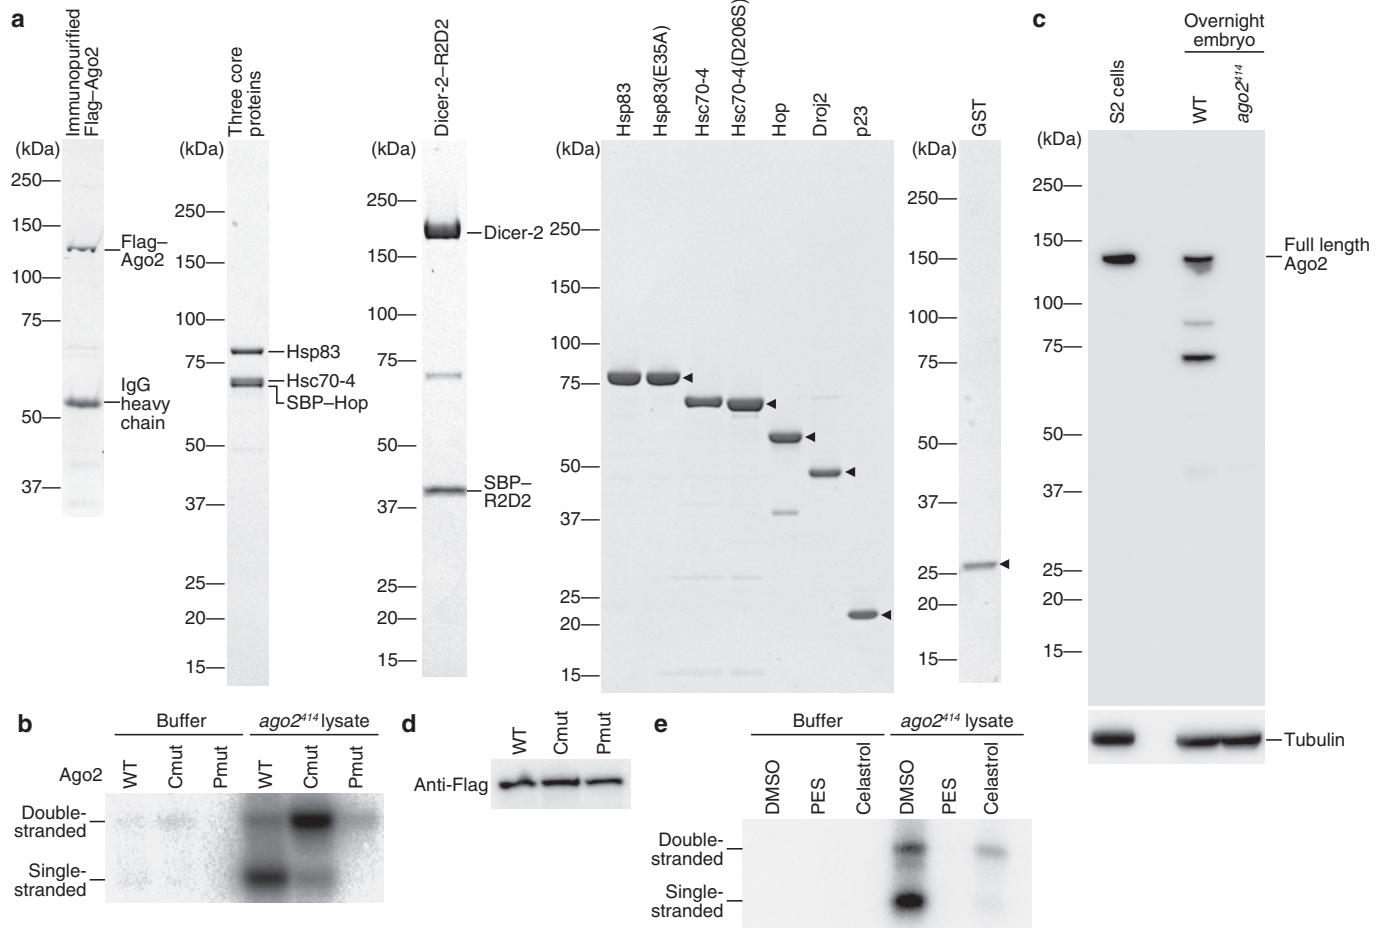
Snapshot observation of wash-resistant spots. For snapshot observation (Figs 2b, c, 4a, b and Extended Data Figs 6e-h and 10d, e), S2 cell lysate containing biotin-labelled Halo-tagged Ago2 was diluted fivefold with lysis buffer just before immobilization to ensure the discreteness of the fluorescent spots. To measure roughly the time course of spot appearance (Extended Data Fig. 6c), the reaction mixture (see earlier) was infused into the observation chamber, and then three random and non-overlapping places were chosen and imaged from one flow chamber at each time point for 1 s at a frame rate of 10 frames s⁻¹ at a power density of 2.5 W mm⁻² for the Ar laser and 360 mW mm⁻² for the He-Ne laser. Images were averaged over 10 frames and deconvoluted, and the total number of spots in the three images was counted for each time point. For snapshot observation of wash-resistant spots, the reaction mixture was infused into the observation chamber and incubated for 1 h. After washing the chamber by wash buffer, the observation mixture was infused into the chamber, and three random and non-overlapping images were taken from each chamber for 1 s as described earlier. For each condition, at least three independent experiments were performed and the total number of spots within 380 × 190 pixels (800 µm²) in the centre of the view field in every three images taken from one flow chamber was counted. The images were averaged, deconvoluted, and collectively analysed as described earlier. It is noted that essentially the same results were obtained regardless of deconvolution.

Snapshot observation after protease treatment. For snapshot observation after protease treatment (Extended Data Fig. 6e, f), the protease mixture containing 1 U µl⁻¹ TurboTEV (T0102S, Accelagen) or 0.1 U µl⁻¹ Turbo3C (H0101S, Accelagen) protease in 20 µl of lysis buffer was infused into the wild-type Ago2-tethered chamber, which had been incubated with the reaction mixture containing all of the factors and imaged, and incubated for 30 min. After incubation, the chamber was washed and imaged as described earlier.

Continuous monitoring of spot appearance. S2 cell lysate containing biotin-labelled Halo-tagged Ago2 was diluted tenfold (Figs 3a-d, 4c-e and Extended Data Figs 7a, 8d and 10f, g) or fivefold (Fig. 3e-h and Extended Data Fig. 7b-e; twice-denser condition to determine the time constant for RISC maturation) with lysis buffer just before immobilization. For the WT+TEV condition (Fig. 3b, d), lysate was incubated with 1 U µl⁻¹ TurboTEV for 30 min in solution and then diluted. Immediately after 20 µl of the reaction mixture (see earlier) was infused into the observation chamber, images were continuously taken for 1,200 s at a frame rate of 1 frames s⁻¹ at a power density of 250 mW mm⁻² for the Ar laser and 36 mW mm⁻² for the He-Ne laser. Stage drift was corrected using a slice alignment plugin⁴². For each condition, at least three independent experiments were performed. To analyse single-molecule association and dissociation events, we first created a maximum intensity projection over all 1,200 frames in stack. Then, we picked the Alexa555 and 647 colocalized fluorescent spots within 380 × 190 pixels (800 µm²) in the centre of the view field of each split image taken from one chamber. In Fig. 3e-h and Extended Data Fig. 7b-e, we analysed the spots that displayed within the entire view field (512 × 256 pixels; 1,450 µm²). We then generated integrated intensity traces to identify concurrent landing events of Alexa555 and 647. Duration times ($t_2 - t_1$) were determined by finding the appearance time of the simultaneous Alexa555 and 647 signals (t_1) and subtracting it from the disappearance time of either or both signals (t_2).

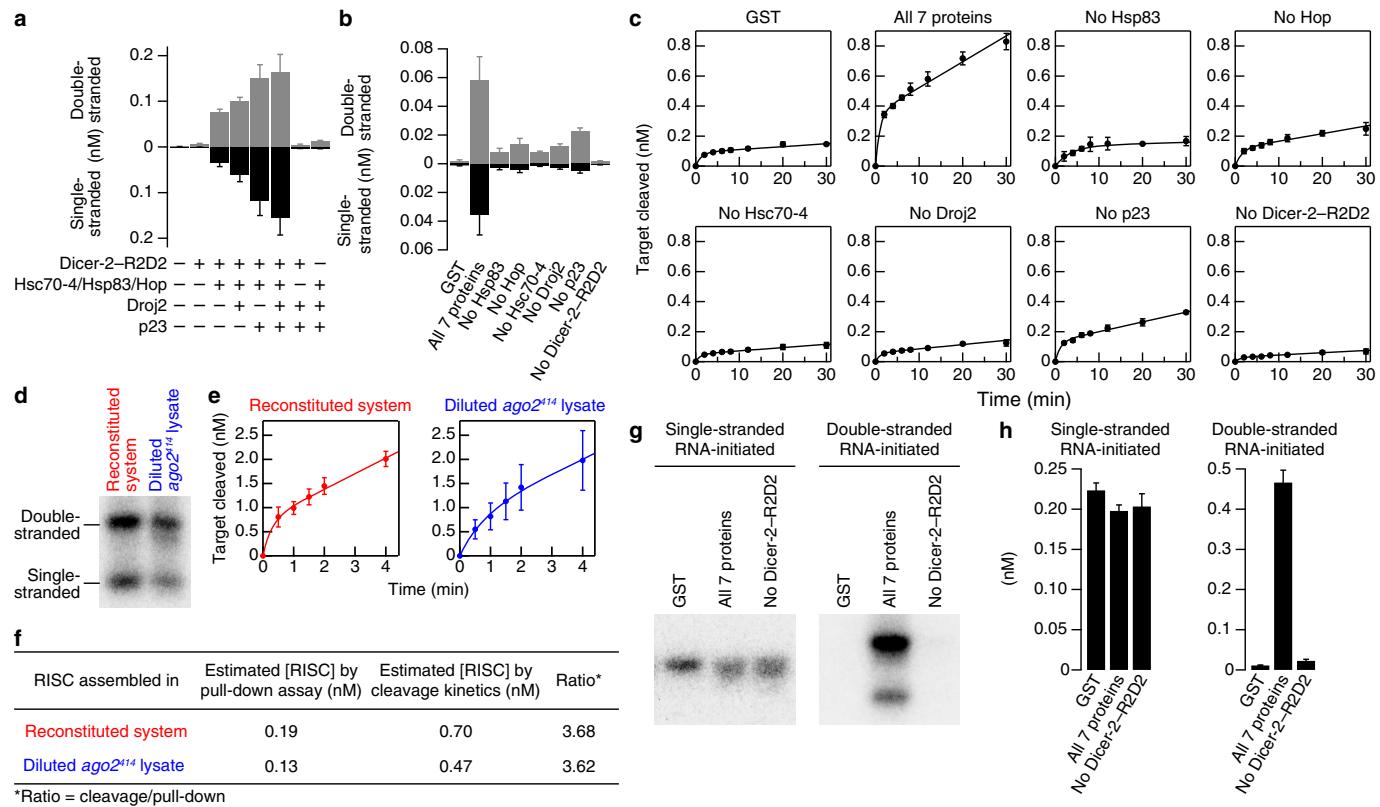
Fluorescence intensity of single-molecule benchmarks. For measuring fluorescence intensity of single-molecule benchmarks (Extended Data Fig. 10e, g), we prepared S2 cell lysate containing biotin-labelled Halo-tagged Ago2(D965A) programmed with fluorescently labelled siRNA duplex, as described earlier. Three kinds of fluorescently labelled duplex (guide/passenger = Alexa647/555, 647/647 and 555/555) were used for programming. The prepared lysate was diluted fivefold with lysis buffer, and was infused into the observation chamber. After washing the chamber by wash buffer, images were taken and analysed under the same acquisition condition and procedure for snapshot observation or continuous monitoring. Three independent chambers were used for imaging for each fluorescently labelled duplex.

31. Nykänen, A., Haley, B. & Zamore, P. D. ATP requirements and small interfering RNA structure in the RNA interference pathway. *Cell* **107**, 309–321 (2001).
32. Haley, B., Tang, G. & Zamore, P. D. *In vitro* analysis of RNA interference in *Drosophila melanogaster*. *Methods* **30**, 330–336 (2003).
33. Förstmann, K. *et al.* Normal microRNA maturation and germ-line stem cell maintenance requires Loquacious, a double-stranded RNA-binding domain protein. *PLoS Biol.* **3**, e236 (2005).
34. Tomari, Y., Du, T. & Zamore, P. D. Sorting of *Drosophila* small silencing RNAs. *Cell* **130**, 299–308 (2007).
35. Kawamata, T. & Tomari, Y. Native gel analysis for RISC assembly. *Methods Mol. Biol.* **725**, 91–105 (2011).
36. Izumi, N. *et al.* AAA+ proteins RUVBL1 and RUVBL2 coordinate PIKK activity and function in nonsense-mediated mRNA decay. *Sci. Signal.* **3**, ra27 (2010).
37. Tsutsumi, A., Kawamata, T., Izumi, N., Seitz, H. & Tomari, Y. Recognition of the pre-miRNA structure by *Drosophila* Dicer-1. *Nature Struct. Mol. Biol.* **18**, 1153–1158 (2011).
38. Haley, B. & Zamore, P. D. Kinetic analysis of the RNAi enzyme complex. *Nature Struct. Mol. Biol.* **11**, 599–606 (2004).
39. Pall, G. S. & Hamilton, A. J. Improved northern blot method for enhanced detection of small RNA. *Nature Protocols* **3**, 1077–1084 (2008).
40. Saito, K. *et al.* Roles for the Yb body components Armitage and Yb in primary piRNA biogenesis in *Drosophila*. *Genes Dev.* **24**, 2493–2498 (2010).
41. Zhou, Z. P. *et al.* Single molecule imaging of the trans-translation entry process via anchoring of the tagged ribosome. *J. Biochem.* **149**, 609–618 (2011).
42. Tseng, Q. *et al.* Spatial organization of the extracellular matrix regulates cell–cell junction positioning. *Proc. Natl Acad. Sci. USA* **109**, 1506–1511 (2012).



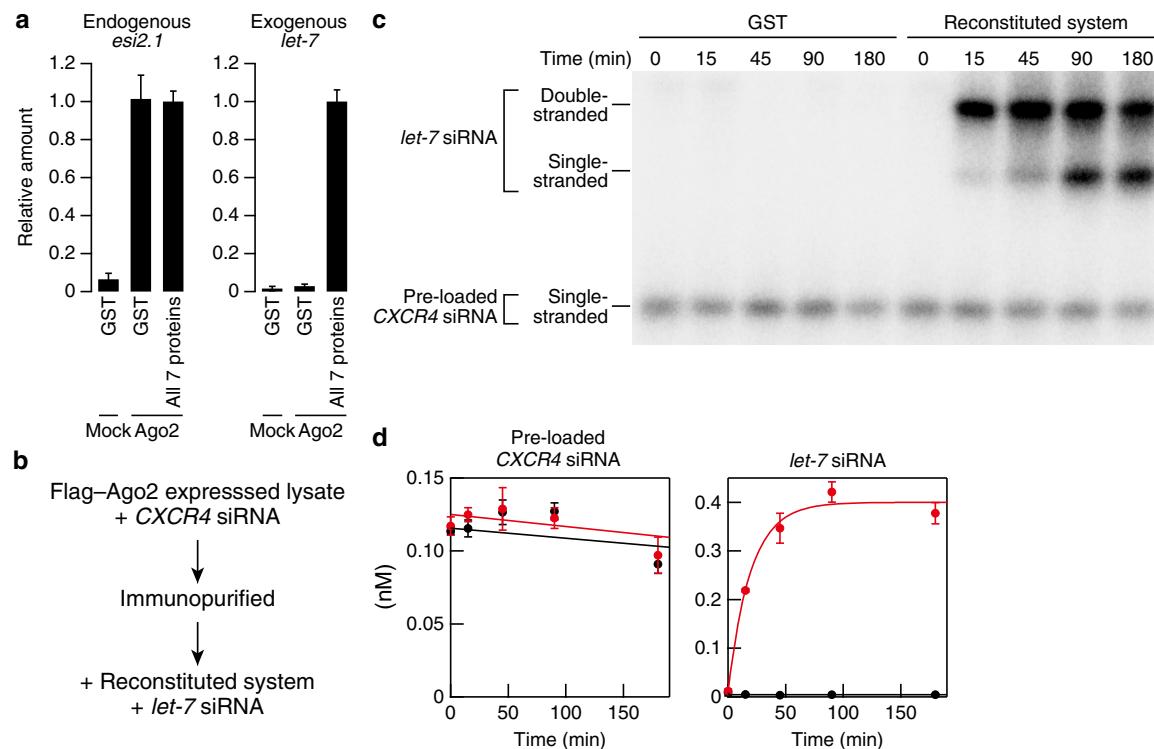
Extended Data Figure 1 | Validation of small RNA pull-down assay. **a**, CBB staining of recombinant proteins used in this study. **b**, Small RNA pull-down assay using immunopurified wild type (WT), catalytic mutant (Cmut; D965A) and 5' nucleotide-binding pocket mutant (Pmut; Y897E) of Ago2. Ago2 alone (Buffer) could not load siRNA, but supplementation of lysate from $ago2^{414}$ mutant flies directed efficient formation of pre- and mature Ago2-RISC. The catalytic mutant Ago2 predominantly formed

pre-Ago2-RISC, while the 5' nucleotide-binding pocket mutant failed to load siRNA. **c**, Western blotting of Ago2 and tubulin in lysate from S2 cells, wild-type embryo and $ago2^{414}$ mutant embryo. **d**, Anti-Flag western blotting of wild-type and mutant Ago2 proteins in lysate. All the constructs are expressed at similar levels. **e**, Hsp70 inhibitor (PES) and p23 inhibitor (Celastrol) block duplex loading into Ago2.

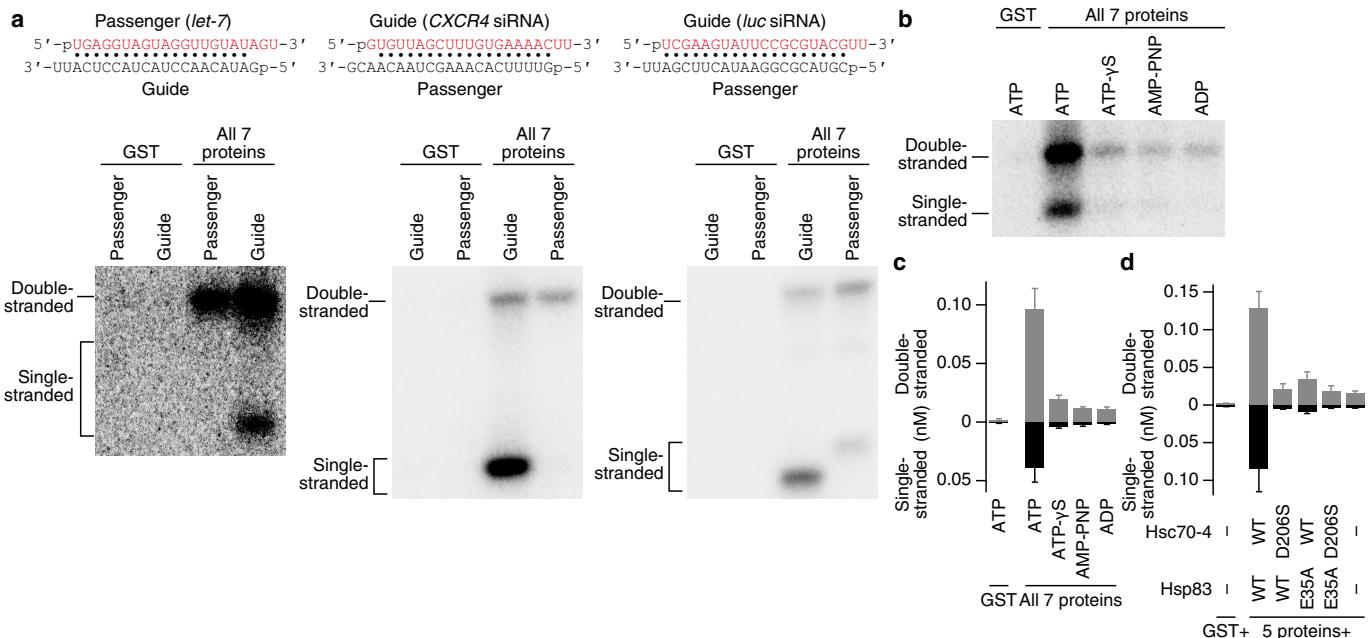


Extended Data Figure 2 | In vitro reconstitution of Ago2-RISC assembly by seven recombinant proteins. **a**, Quantification of Fig. 1a. **b**, Quantification of Fig. 1b. **c**, Time course of target cleavage in the reconstituted system. Figure 1c represents the data at 30 min. **d**, Activities of the reconstituted system and *ago2*⁴¹⁴ mutant lysate. **e**, Cleavage assay in a multiple turnover condition by Ago2-RISCs assembled in the reconstituted system or *ago2*⁴¹⁴ mutant embryo lysate (**d**). *ago2*⁴¹⁴ lysate was diluted to normalize the amount of the assembled Ago2-RISC. **f**, Estimated concentrations of mature single-strand-containing Ago2-RISCs by the pull-down assay (**d**) and the pre-steady-state

'burst' kinetics of target cleavage (**e**). Owing to a technical limitation in the quantification of small amounts of radiolabelled RNAs, the concentrations of RISCs estimated by the burst kinetics were higher than those estimated by the pull-down assay. However, the ratios between them were essentially identical between RISCs assembled in the reconstituted system and *ago2*⁴¹⁴ mutant lysate, suggesting that the reconstituted Ago2-RISC is as fully active as the one assembled in lysate. **g**, Dicer-2-R2D2 and the chaperone machinery are dispensable for loading of single-stranded RNAs whereas they are essential for duplex loading. **h**, Quantification of **g**. Data represent mean \pm s.d. ($n = 3$).

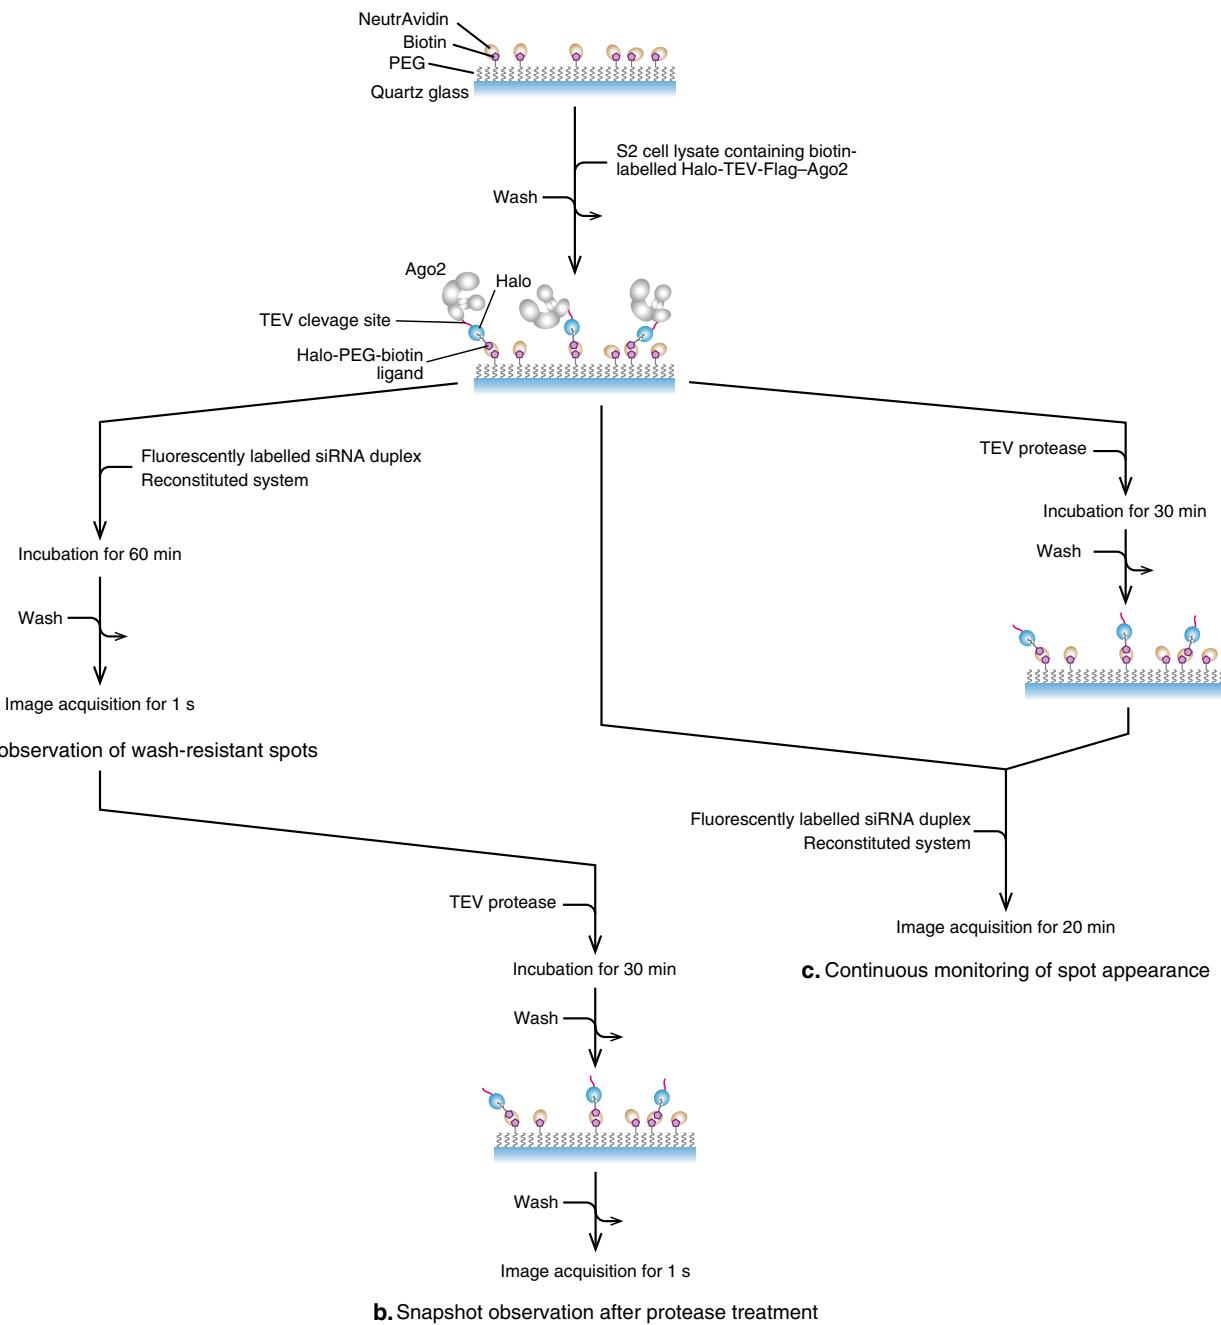


Extended Data Figure 3 | Unloading is not accompanied by loading. **a**, Quantification of Fig. 1d. **b**, Experimental strategy for c. **c**, The amount of pre-loaded CXCR4 siRNA did not change during new loading of *let-7* siRNA. **d**, Quantification of c. Red, reconstituted system; black, GST. Data represent mean \pm s.d. ($n = 3$).



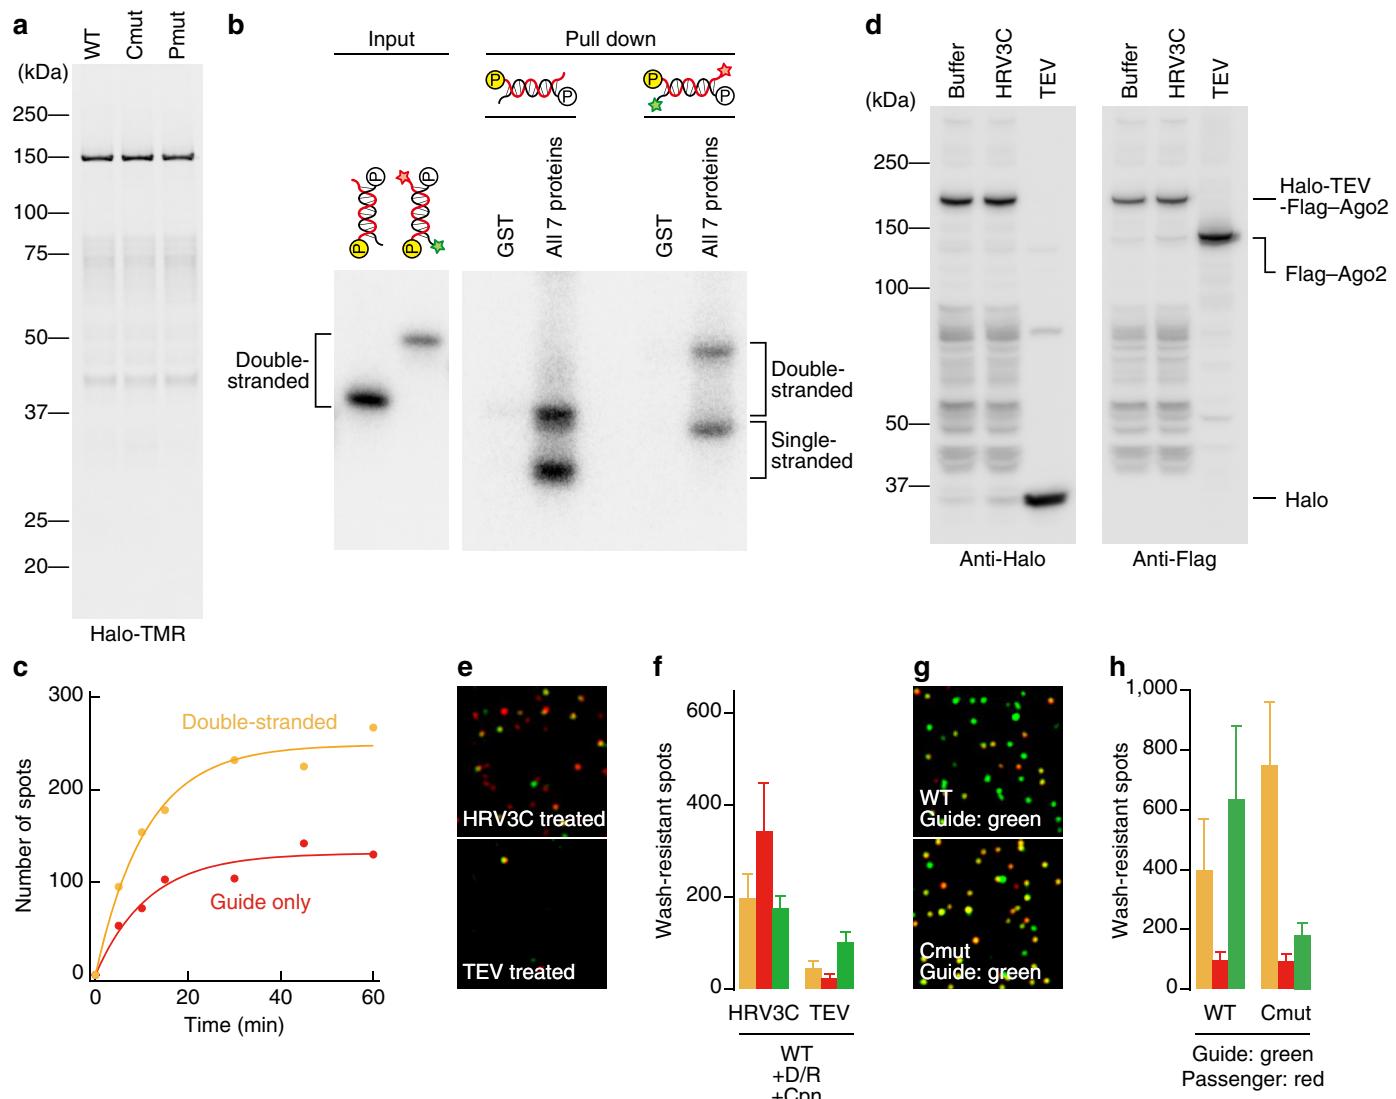
Extended Data Figure 4 | The reconstituted system faithfully recapitulates siRNA loading asymmetry and ATP dependency. **a**, Ago2-RISC was assembled with siRNAs with either the guide strand or the passenger strand radiolabelled. Left, asymmetry-flipped *let-7* siRNA. In contrast to the normal *let-7* siRNA (Fig. 1e), the *let-7* sequence in this siRNA duplex serves as the passenger strand due to the introduction of a mismatch at the opposite end.

Middle, highly asymmetric CXCR4 siRNA. Right, relatively symmetric *luc* siRNA, in which the two strands can serve as the guide strand with similar chances. **b**, Small RNA pull-down assay was performed in the reconstituted system with 0.5 mM ATP, ATP- γ S, AMP-PNP or ADP. **c**, Quantification of **b**. **d**, Quantification of Fig. 1f.



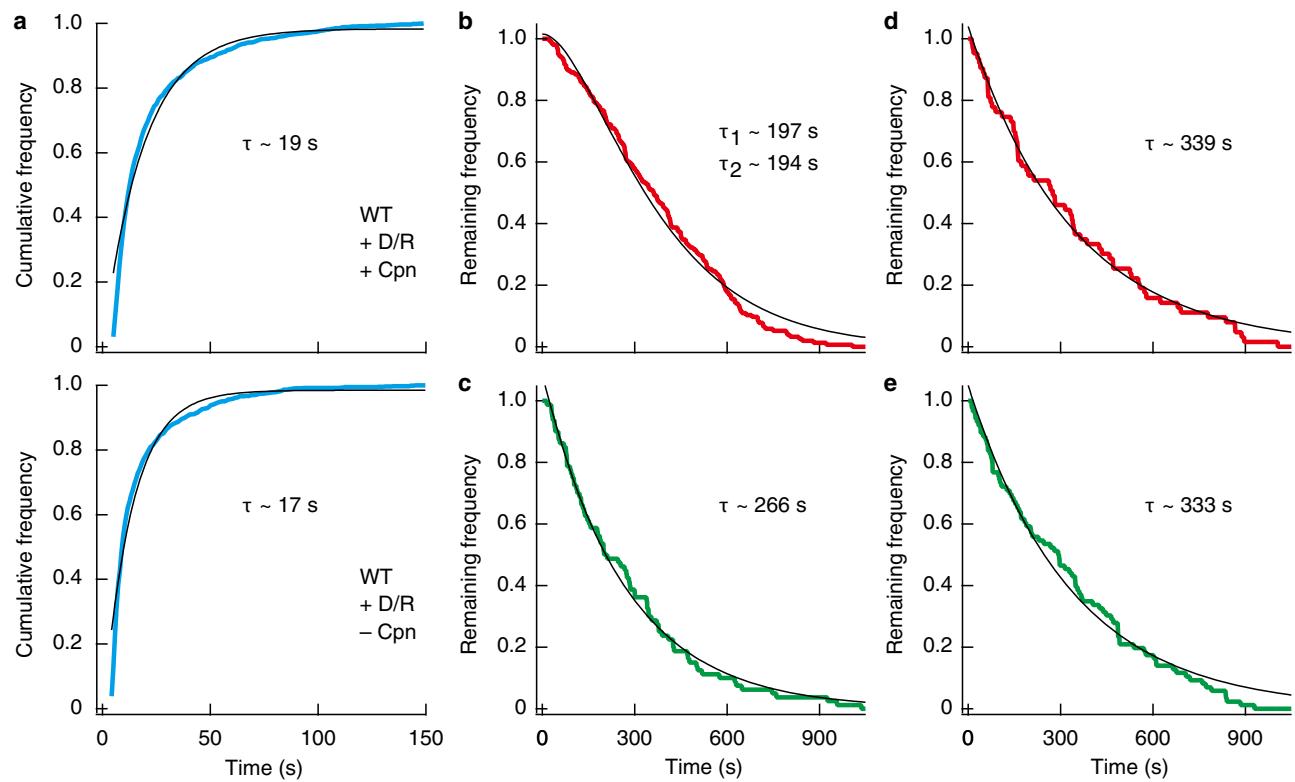
Extended Data Figure 5 | Experimental scheme for single-molecule analysis. a–c, Halo-TEV-Flag-Ago2 was tethered to the surface as illustrated. For snapshot observation (a), the reaction mixture was infused into the observation chamber and incubated for 1 h. After washing of the surface, the images of the wash-resistant spots were acquired for 1 s. For snapshot

observation after protease treatment (b), the protease mixture was infused into the chamber, which had been imaged and incubated for 30 min. For continuous monitoring (c), images were acquired for 20 min without washing immediately after the reaction was started on the surface.



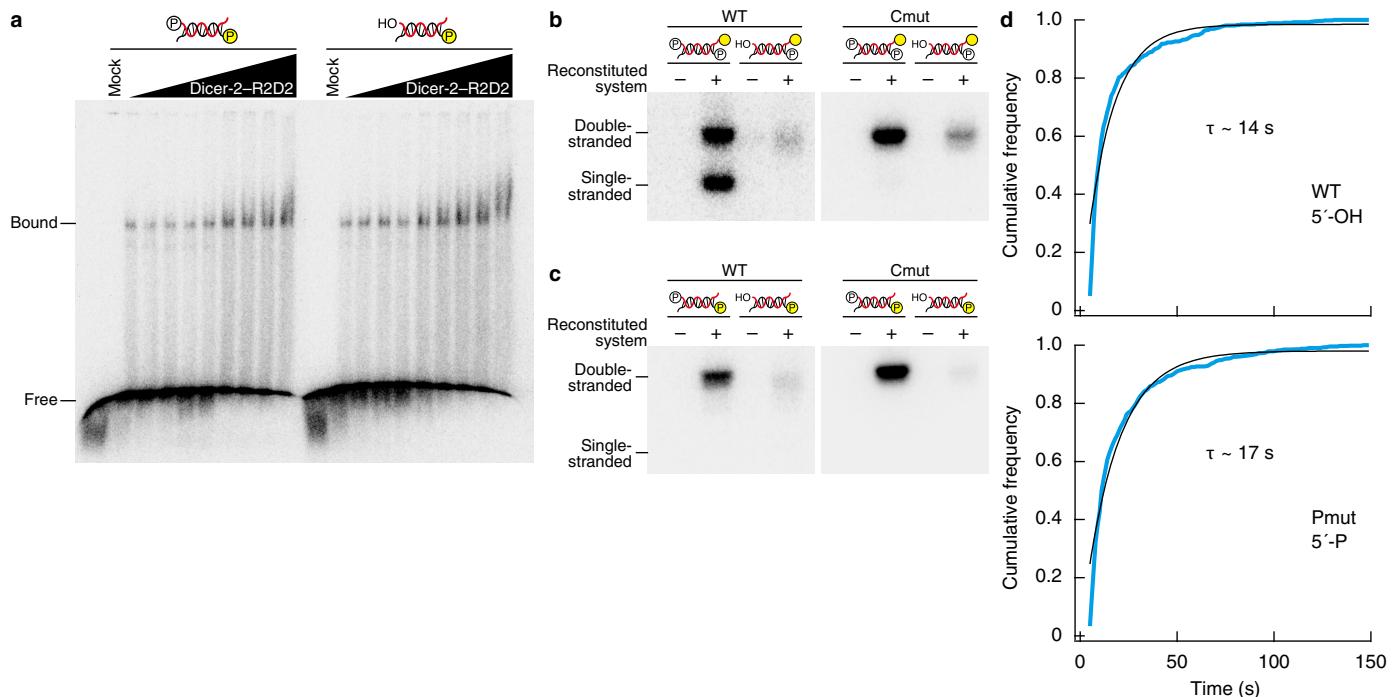
Extended Data Figure 6 | Halo-tagged Ago proteins and fluorescently labelled siRNAs used in single-molecule analysis. **a**, Visualization of Halo-TEV-Flag-tagged wild type, catalytic mutant (Cmut; D965A) and 5' nucleotide-binding pocket mutant (Pmut; Y897E) of Ago2 by a Halo-TMR ligand. **b**, RNA pull-down assay with siRNAs with or without 3' fluorescent modifications. Red, guide strand; black, passenger strand. The 5'-phosphate of the guide strand was radiolabelled as denoted by yellow circles. Fluorescent labelling did not compromise the efficiency of RISC assembly *in vitro*. **c**, Time course of the guide-alone (red) or double-stranded (yellow) spot formation for wild-type Ago2 with the reconstituted system. **d**, TEV-protease-mediated cleavage of Halo-TEV-Flag-Ago2. After S2 lysate containing Halo-TEV-Flag-Ago2 was incubated with proteases (HRV3C or TEV) or lysis buffer for 30 min,

western blotting was performed using anti-Halo and anti-Flag antibodies. **e, f**, Representative single-molecule images (**e**) and quantification (**f**) after protease treatment for 30 min (Extended Data Fig. 5b). Wash-resistant spots observed in the WT, +D/R and +Cpn conditions (Fig. 2b, c) largely disappeared after TEV-protease-mediated cleavage but not after the control treatment of HRV3C protease. The colour scheme is the same as in Fig. 2b, c. **g, h**, Validation of pre- and mature RISC formation with inversion of the strand colours. Representative single-molecule images (**g**) and quantification (**h**) of wash-resistant spots after 60 min incubation with the reconstituted system together with an siRNA with a 3' Alexa555 (green)-labelled guide strand and a 3' Alexa647 (red)-labelled passenger strand. Data represent mean \pm s.d. (WT in **e**, $n = 4$; others, $n = 3$).



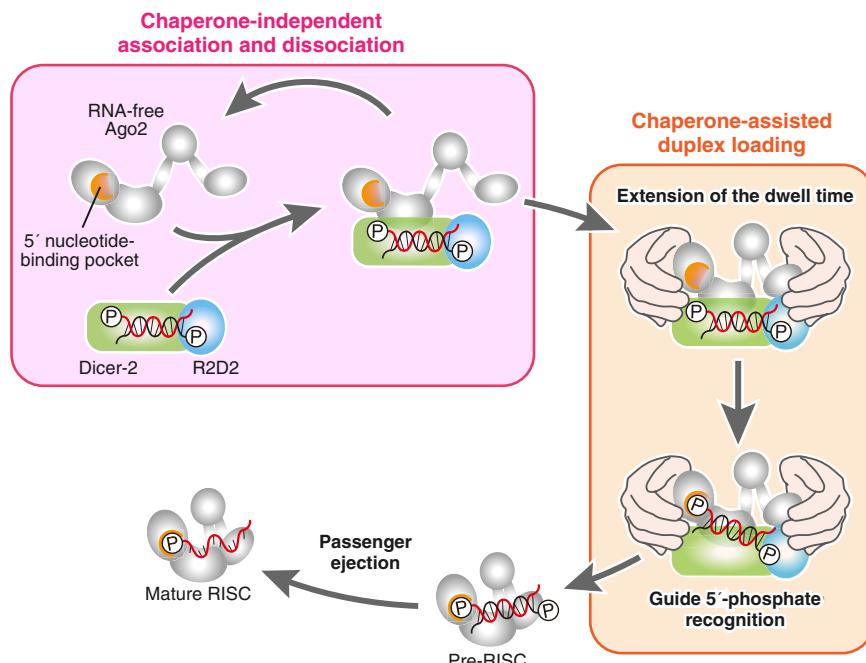
Extended Data Figure 7 | Determination of the time constants by cumulative curve fitting. **a**, The data in Fig. 3c for short binding were accumulated and fitted to the following equation derived from the one-step reaction model: $C \cdot [1 - \exp(-t/\tau)]$, where C is the normalized parameter, τ is the time constant, and t is the reaction time. **b**, The data in Fig. 3e for the yellow-to-red transition in wild type were accumulated and fitted to

the following equation derived from the two-step reaction model: $[C_0/(C_1 - C_2)] \cdot [C_1 \cdot \exp(-C_2 \cdot t) - C_2 \cdot \exp(-C_1 \cdot t)]$, where C_0 is the normalized parameter, C_1 and C_2 are the inverse of the time constants τ_1 and τ_2 , respectively, and t is the reaction time. **c–e**, The data in Fig. 3f–h were accumulated and fitted to the following equation: $C \cdot \exp(-t/\tau)$, where C is the normalized parameter, τ is the time constant, and t is the reaction time.



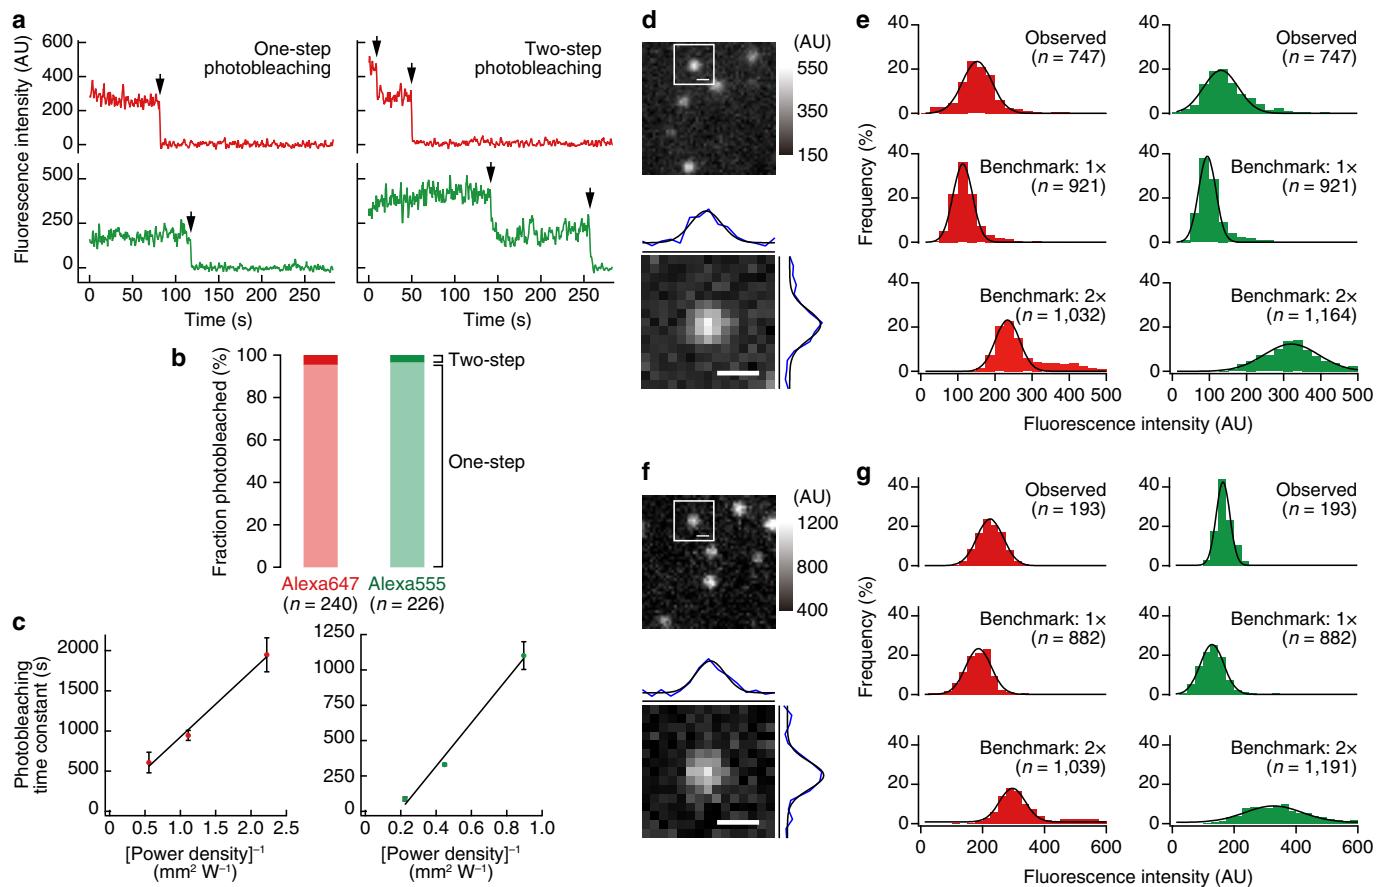
Extended Data Figure 8 | The 5'-phosphate of the guide strand is dispensable for binding to Dicer-2-R2D2 but is essential for duplex loading into Ago2. **a**, Gel shift assay with increasing concentrations of recombinant Dicer-2-R2D2 and passenger-strand-radiolabelled siRNA duplex with the 5'-phosphate (left) or 5'-hydroxyl group (right) on the guide strand. The 5'-phosphate of the guide strand is dispensable for the interaction with Dicer-2-R2D2. Red, guide strand; black, passenger strand. **b, c**, Small RNA pull-down assay with siRNA duplexes containing 5'-phosphate (left) or 5'-hydroxyl group (right) on the guide strand with wild-type or the catalytic

mutant Ago2. The 3' end of the guide strand (**b**) or the 5' end of the passenger strand (**c**) of the siRNA duplex was radiolabelled, as denoted by yellow circles. The 5'-phosphate on the guide strand is essential for duplex loading into Ago2; in the absence of the guide 5'-phosphate, the siRNA duplex is hardly loaded while maintaining the functional asymmetry. **d**, Determination of the time constants of short binding in the absence of 5'-phosphate recognition. The data in Fig. 4d were accumulated and fitted to the one-step reaction model as in Extended Data Fig. 7a.



Extended Data Figure 9 | Multiple fundamental steps in the assembly of *Drosophila* Ago2-RISC. In the absence of the chaperone machinery, an siRNA bound to Dicer-2-R2D2 binds Ago2 only transiently, dissociating rapidly. The Hsc70/Hsp90 chaperone machinery extends their dwell time, presumably by inducing a productive state of Ago2, in a manner dependent on

the recognition of the 5'-phosphate on the guide strand by Ago2. Successful anchoring of the 5'-phosphate would trigger loading of the rest of the siRNA duplex from Dicer-2-R2D2 into Ago2, leading to passenger ejection and mature RISC formation.



Extended Data Figure 10 | Single-molecule properties of the observed fluorescent spots. **a**, Representative time courses of one-step (left) and two-step (right) photobleaching for the guide (red) or passenger (green) strand. Spots for wild-type Ago2 with the reconstituted system were analysed. Arrows show the timings of photobleaching. **b**, ~95% of the fluorescent spots showed one-step photobleaching for each colour, indicating that each spot contained a single RNA molecule. **c**, Photobleaching dependence on power density. Left, Alexa647. Right, Alexa555. The photobleaching time constants of Alexa647 and 555 under continuous monitoring condition were estimated as $\sim 22,800$ and $\sim 5,900$ s, respectively. Data represent mean \pm s.d. ($n = 3$). **d**, Representative observed spots in snapshot observation without deconvolution (raw data). The distribution of fluorescence intensity was fitted to Gaussian function (bottom). The diameter of the spots is comparable to the diameter of the Airy disk (533 nm), indicating that the spots are diffraction limited. Scale bar, 500 nm. **e**, The signal intensity distributions of observed spots and single-molecule benchmarks for Alexa647 (red, left) and Alexa555 (green, right) in snapshot observation. Observed, fluorescent spots observed in snapshot observation (Fig. 2c; WT, +D/R, +Cpn, yellow). Benchmark, the catalytic mutant Ago2 programmed with siRNA duplex, of which the guide and passenger strands were labelled with Alexa647 and 555, respectively (1 \times , middle) or with the same dye (2 \times , bottom). Images of single-molecule benchmarks were taken under the same acquisition condition of snapshot observation. Background was subtracted from the raw data. **f**, Representative observed spots in continuous monitoring without deconvolution (raw data). The distribution of fluorescence intensity was fitted to Gaussian function (bottom). The diameter of the spots is comparable to the diameter of the Airy disk (533 nm), indicating that the spots are diffraction limited. Scale bar, 500 nm. **g**, The signal intensity distributions of observed spots and single-molecule benchmarks for Alexa647 (red, left) and Alexa555 (green, right) in continuous monitoring. Observed, fluorescent spots observed in continuous monitoring (Fig. 3d; WT, +D/R, +Cpn, yellow). Benchmark, the catalytic mutant Ago2 programmed with siRNA duplex, of which the guide and passenger strands were labelled with Alexa647 and 555, respectively (1 \times , middle) or with the same dye (2 \times , bottom). Images of single-molecule benchmarks were taken under the same acquisition condition of continuous monitoring. Background was subtracted from the raw data.

catalytic mutant Ago2 programmed with siRNA duplex, of which the guide and passenger strands were labelled with Alexa647 and 555, respectively (1 \times , middle) or with the same dye (2 \times , bottom). Images of single-molecule benchmarks were taken under the same acquisition condition of snapshot observation. Background was subtracted from the raw data. **f**, Representative observed spots in continuous monitoring without deconvolution (raw data). The distribution of fluorescence intensity was fitted to Gaussian function (bottom). The diameter of the spots is comparable to the diameter of the Airy disk (533 nm), indicating that the spots are diffraction limited. Scale bar, 500 nm. **g**, The signal intensity distributions of observed spots and single-molecule benchmarks for Alexa647 (red, left) and Alexa555 (green, right) in continuous monitoring. Observed, fluorescent spots observed in continuous monitoring (Fig. 3d; WT, +D/R, +Cpn, yellow). Benchmark, the catalytic mutant Ago2 programmed with siRNA duplex, of which the guide and passenger strands were labelled with Alexa647 and 555, respectively (1 \times , middle) or with the same dye (2 \times , bottom). Images of single-molecule benchmarks were taken under the same acquisition condition of continuous monitoring. Background was subtracted from the raw data.

MAD2L2 controls DNA repair at telomeres and DNA breaks by inhibiting 5' end resection

Vera Boersma^{1*}, Nathalie Moatti^{1*}, Sandra Segura-Bayona¹, Marieke H. Peuscher¹, Jaco van der Torre¹, Brigitte A. Wevers¹, Alexandre Orthwein², Daniel Durocher^{2,3} & Jacqueline J. L. Jacobs¹

Appropriate repair of DNA lesions and the inhibition of DNA repair activities at telomeres are crucial to prevent genomic instability. By fuelling the generation of genetic alterations and by compromising cell viability, genomic instability is a driving force in cancer and ageing^{1,2}. Here we identify MAD2L2 (also known as MAD2B or REV7) through functional genetic screening as a novel factor controlling DNA repair activities at mammalian telomeres. We show that MAD2L2 accumulates at uncapped telomeres and promotes non-homologous end-joining (NHEJ)-mediated fusion of deprotected chromosome ends and genomic instability. MAD2L2 depletion causes elongated 3' telomeric overhangs, indicating that MAD2L2 inhibits 5' end resection. End resection blocks NHEJ while committing to homology-directed repair, and is under the control of 53BP1, RIF1 and PTIP³. Consistent with MAD2L2 promoting NHEJ-mediated telomere fusion by inhibiting 5' end resection, knockdown of the nucleases CTIP or EXO1 partially restores telomere-driven genomic instability in MAD2L2-depleted cells. Control of DNA repair by MAD2L2 is not limited to telomeres. MAD2L2 also accumulates and inhibits end resection at irradiation-induced DNA double-strand breaks and promotes end-joining of DNA double-strand breaks in several settings, including during immunoglobulin class switch recombination. These activities of MAD2L2 depend on ATM kinase activity, RNF8, RNF168, 53BP1 and RIF1, but not on PTIP, REV1 and REV3, the latter two acting with MAD2L2 in translesion synthesis⁴. Together, our data establish MAD2L2 as a crucial contributor to the control of DNA repair activity by 53BP1 that promotes NHEJ by inhibiting 5' end resection downstream of RIF1.

As the processes underlying telomere-driven genomic instability are not completely understood, we performed a functional genetic screen to identify telomere-induced genomic instability regulators (TIGIRs). The TIGIR screen relies on well-controlled reversible inactivation of telomere component TRF2 by expressing temperature-sensitive TRF2(Ile468Ala) (TRF2ts) in *Trf2*^{-/-} *p53*^{-/-} (also known as *Trf2*^{-/-} *Trp53*^{-/-}) mouse embryo fibroblasts (MEFs)⁵. At the permissive temperature (32 °C), the TRF2ts MEFs have intact TRF2-mediated telomere protection, but at non-permissive temperatures (37–39 °C) TRF2ts is inactive, causing ATM kinase activation, accumulation of DNA damage response (DDR) proteins and NHEJ-dependent ligation at chromosome ends⁵. In the TIGIR screen (Fig. 1a), prolonged TRF2-inactivation causes TRF2ts cells to irreversibly arrest or die owing to severe chromosome end fusion that drives cells into genomic crisis^{6–8}. However, diminished telomere fusion, such as with RNF8 or DNA ligase IV deficiency, allows survival and proliferation despite telomere uncapping^{9–11}.

In a triplicate TIGIR screen we assayed 1,976 short hairpin RNAs (shRNAs), targeting 391 genes linked to DDR, for shRNAs that allow TRF2ts cells to survive 12 days of telomere uncapping. Among shRNA targets substantially enriched in at least two out of three screens were several factors previously shown to control telomere fusion (Extended Data Fig. 1a, b), including ATM, NBS1, RAD50, 53BP1 and RNF8

(refs 11–15). Remarkably, the most prominent screen hit, also independently recovered in a genome-wide TIGIR screen (our unpublished results), was MAD2L2 (Extended Data Fig. 1a, b). MAD2L2 has no known function at telomeres but acts as a non-catalytic interaction partner of REV1 and REV3 in translesion synthesis⁴. Knockdown of *Mad2l2* with independent shRNAs markedly increased the survival of TRF2ts MEFs subjected to telomere uncapping, which was abolished by complementation with exogenous shRNA-resistant human *MAD2L2* (Fig. 1b, c and Extended Data Fig. 1c–f). Interestingly, we failed to detect a similar activity for REV1 or REV3 (Extended Data Fig. 1g–i).

Indeed, enhanced survival of MAD2L2-depleted cells after long-term TRF2 inactivation was due to diminished telomere fusion. *Mad2l2* knockdown significantly reduced telomere fusions, which was prevented by complementation with exogenous *MAD2L2* (Fig. 2a and Extended Data Fig. 2a). Telomeres end with a single-stranded 3' G-overhang that undergoes DNA-ligase-IV/NHEJ-dependent degradation after loss of TRF2-mediated telomere protection^{9,10}. Analysis of telomeric 3' G-overhangs showed that MAD2L2 depletion prevents overhang loss at 48 h of TRF2 inactivation, further confirming that MAD2L2 is essential for efficient processing of deprotected telomeres by NHEJ (Fig. 2b and Extended Data Fig. 2b). The defect in NHEJ-mediated telomere fusion in *Mad2l2*-knockdown cells was not explained by changes in cell-cycle progression (Fig. 1b, 32 °C panels, and Extended Data Figs 1e and 2c, d). Moreover, MAD2L2 is essential for sister-telomere fusion after activation of DNA repair in mitosis¹⁶ (Extended Data Fig. 2e). In line

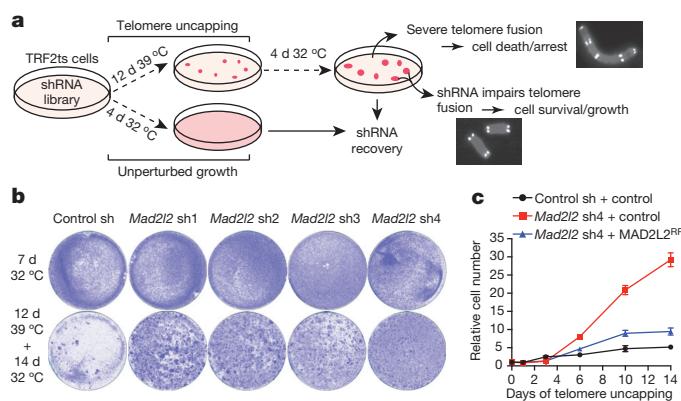


Figure 1 | A functional genetic screen identifies MAD2L2 as a critical factor in telomere-driven genomic instability. **a**, Outline of TIGIR screen to identify factors controlling telomere-driven genomic instability. **b**, Survival assays of TRF2ts MEFs infected with control shRNA (sh) or *Mad2l2* shRNAs (sh1–sh4), stained after growth as indicated. **c**, Growth curves at 39 °C of TRF2ts cells transduced with control or *Mad2l2* sh4 shRNAs and complemented with empty control or RNA interference (RNAi)-resistant Flag-MAD2L2^{RR} (data are mean ± s.d. from quadruplicate technical replicates).

¹Division of Molecular Oncology, The Netherlands Cancer Institute, Plesmanlaan 121, 1066 CX, Amsterdam, The Netherlands. ²The Lunenfeld-Tanenbaum Research Institute, Mount Sinai Hospital, 600 University Avenue, Toronto, Ontario M5G 1X5, Canada. ³Department of Molecular Genetics, University of Toronto, Ontario M5S 3E1, Canada.

*These authors contributed equally to this work.

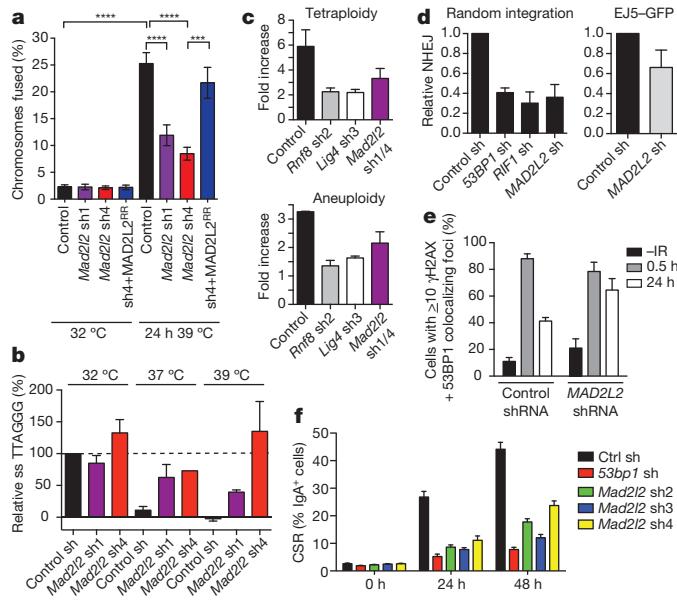
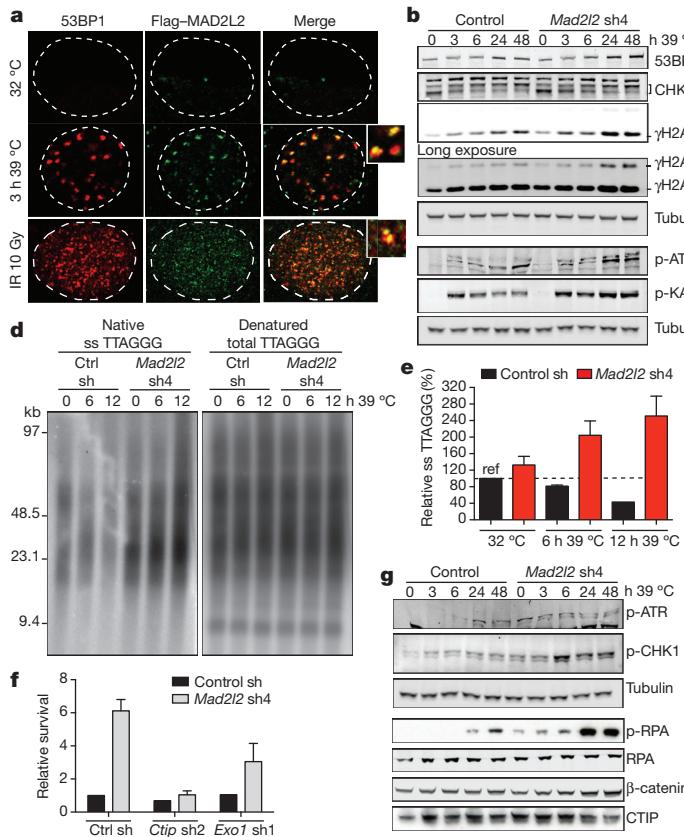


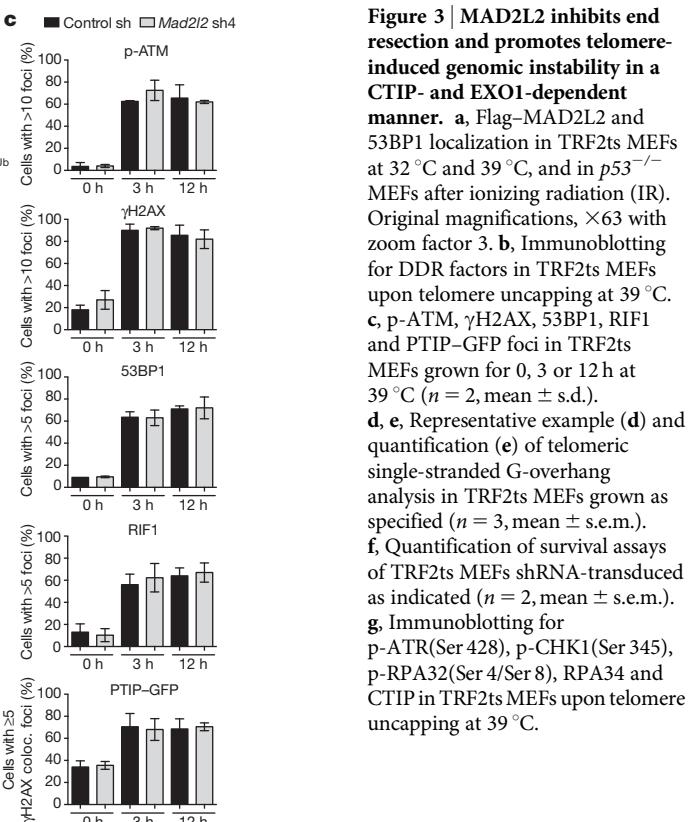
Figure 2 | MAD2L2 facilitates telomeric G-overhang degradation, NHEJ-mediated telomere fusion, repair and CSR. **a**, Chromosome fusion in shRNA-transduced TRF2ts MEFs ($n = 2–4$, mean \pm s.e.m., *** $P = 0.0002$, **** $P < 0.0001$, Kruskal–Wallis analysis of variance (ANOVA)). **b**, Telomeric single-stranded (ss) G-overhang quantification in TRF2ts MEFs at 32 °C and after 48 h at 37 °C or 39 °C ($n = 4$, except *Mad2l2* sh4 at 37 °C: $n = 1$, mean \pm s.e.m.). **c**, Tetraploidy (top) and aneuploidy (bottom) after telomere uncapping in TRF2ts MEFs ($n = 2$, mean \pm s.d.). **d**, NHEJ-mediated repair in U2OS cells analysed by random plasmid integration (left, $n = 3$) or an EJ5-GFP reporter containing green fluorescent protein (GFP) (right, $n = 6$, mean \pm s.d.). **e**, MAD2L2 depletion impairs resolution of DSBs after irradiation (2 Gy) in HDR-deficient (RAD51-depleted) U2OS cells, indicating defective NHEJ ($n = 2$, mean \pm s.e.m.). **f**, CSR in shRNA-transduced CH12F3-2 cells ($n = 3$, mean \pm s.d.). Ctrl, control.



with MAD2L2 promoting telomere fusion, chromosomal rearrangement and mis-segregation^{6–8}, MAD2L2 depletion prevented tetraploidy and aneuploidy after telomere uncapping (Fig. 2c and Extended Data Fig. 2f).

To address whether the requirement of MAD2L2 for NHEJ extends beyond telomeres, we investigated whether MAD2L2 depletion affects end-joining in other settings. Knockdown of human *MAD2L2* considerably impaired NHEJ-mediated repair of I-Sce1-induced DNA double-strand breaks (DSBs), NHEJ-mediated random plasmid integration, resolution of irradiation-induced DSBs in homology-directed repair (HDR)-deficient cells and clonogenic survival after irradiation, similar to depletion of the NHEJ-promoting factors 53BP1 and RIF1 (Fig. 2d, e and Extended Data Fig. 3a–e). Moreover, class switch recombination (CSR), a physiological process depending on joining of DNA DSBs by NHEJ¹⁷, was substantially impaired in MAD2L2-depleted cells (Fig. 2f and Extended Data Fig. 3f–i). This CSR defect was not due to defects in cell proliferation or expression of *Aid* (also known as *Aicda*) messenger RNA, IgM and IgA germ-line transcripts. Thus, MAD2L2 promotes NHEJ-mediated repair in several contexts.

To investigate how MAD2L2 promotes NHEJ, we first addressed its localization. MAD2L2 accumulated both at TRF2-deprotected telomeres and irradiation-induced DSBs (Fig. 3a), implying a direct activity of MAD2L2 at these sites. Next we examined activation of the DDR. At 3, 6 and 12 h of TRF2 inactivation, both control and MAD2L2-depleted cells activated ATM, phosphorylated H2AX, KAP1 and CHK2 (γH2AX, p-KAP1 and p-CHK2, respectively), and formed p-ATM and γH2AX foci to a similar extent (Fig. 3b, c and Extended Data Fig. 4a, b). Moreover, accumulation of 53BP1, RIF1 and PTIP to uncapped telomeres, as well as irradiation-induced DSBs, was unperturbed in MAD2L2-depleted cells (Fig. 3c and Extended Data Fig. 4a–c). Thus, MAD2L2 is not required for recognition of uncapped telomeres or DSBs as damage, nor for initiation of DDR signalling. However, probably reflecting aberrant processing of deprotected telomeres, we noticed increased DDR signalling at 24 and 48 h of TRF2 inactivation in MAD2L2-depleted cells,



visible as increased H2AX and KAP1 phosphorylation, and H2AX ubiquitylation (Fig. 3b).

Next we examined the terminal structure of telomeres in more detail. Analysis of 3' single-stranded G-overhangs at 6 and 12 h of TRF2 inactivation revealed that in the absence of MAD2L2, telomeres gain considerably longer 3' single-stranded overhangs. At 12 h, *Mad2l2*-knockdown cells contained 2.5-fold higher, 3'-exonuclease sensitive overhang signals than control cells without uncapping, and nearly sixfold higher overhang signals than control cells at 12 h of telomere uncapping (Fig. 3d, e and Extended Data Fig. 5a). Overhang increase upon MAD2L2 depletion was observed in both wild-type and DNA-ligase-IV-deficient TRF2ts cells (Extended Data Fig. 5b-d), indicating it is not a passive consequence of lack of NHEJ-mediated overhang degradation, but represents an active role of MAD2L2 in inhibiting 5' end resection. 5' DNA end resection is a powerful obstruction to NHEJ by creating DNA ends unsuitable for ligation¹⁸, explaining why MAD2L2 depletion impairs NHEJ-mediated repair.

5' end resection involves the combined activities of CTIP, MRN, EXO1, BLM and DNA2, and is blocked by 53BP1 and its interaction partners RIF1 and PTIP^{3,13,19-24}. In line with MAD2L2 promoting NHEJ by impeding 5' end resection, co-depletion of CTIP or EXO1 with MAD2L2 partially reversed the increased survival of *Mad2l2*-knockdown cells undergoing prolonged telomere uncapping (Fig. 3f and Extended Data Fig. 5e-h). Moreover, consistent with the generation of longer single-stranded overhangs that block NHEJ and prime for HDR, *Mad2l2*-knockdown cells showed increased phosphorylation of the ATR target CHK1 and of the single-stranded DNA-binding protein RPA (Fig. 3g). In addition, MAD2L2-depleted cells showed increased HDR at telomeres by increased telomere sister-chromatid exchanges (TSCE; Extended Data Fig. 6a).

Because the effects of MAD2L2 depletion on end resection, telomere fusion and CSR resemble those of 53BP1, RIF1 and, to a certain extent, PTIP inactivation^{3,13,19-24}, we asked whether MAD2L2 is epistatic with

53BP1, RIF1 or PTIP. Similarly, we addressed the requirement of REV1 and REV3 for the ability of MAD2L2 to control end resection and telomere fusion. Unlike in wild-type or PTIP-deficient MEFs, MAD2L2 depletion did not reduce fusion of TRF2-deprotected telomeres in 53BP1- or RIF1-deficient MEFs (Fig. 4a and Extended Data Fig. 6b-e). However, MAD2L2(Leu186Ala) and MAD2L2(Cys70Arg), defective in REV1 or REV3 interaction, respectively^{25,26}, were still capable of promoting telomere fusion and telomere-driven genomic instability (Fig. 4b and Extended Data Fig. 7a, g, i). Moreover, MAD2L2 depletion resulted in increased RPA phosphorylation after irradiation in wild-type, *Ptip*^{-/-}, *Rev1*^{-/-} or *Rev3*^{-/-} MEFs, but not in *53bp1*^{-/-} (also known as *Trp53bp1*^{-/-}) or *Rif1*^{-/-} MEFs (Fig. 4c). This indicates that MAD2L2 controls end resection and fusion of deprotected telomeres in a 53BP1- and RIF1-dependent, but PTIP-, REV1- and REV3-independent manner. Unanticipated, MAD2L2-depleted *Rif1*^{-/-} MEFs showed an increased background of chromosome end fusion caused by telomere association between sister chromatids (Fig. 4a and Extended Data Fig. 6d, e). The molecular nature of these sister telomere associations remains to be determined, but they occur irrespective of TRF2 inhibition and potentially reflect replication problems²⁷.

To address the dependencies for MAD2L2 function in DNA repair further, we investigated the requirements for MAD2L2 accumulation to DSBs and NHEJ activity at telomeres. Both DSB accumulation and the ability to promote telomere-driven genomic instability by MAD2L2 require an intact HORMA domain and carboxy terminus, while potential S/TQ (Ser or Thr residues that precede Gln) phosphorylation on Thr 103 appears dispensable (Fig. 4b and Extended Data Fig. 7). Interestingly, while MAD2L2 accumulation to irradiation-induced DDR foci (IRIFs) was independent of PTIP or interaction of MAD2L2 with REV1 or REV3, it clearly depended on ATM kinase activity, RNF8, RNF168, 53BP1 and RIF1 (Fig. 4d, e and Extended Data Figs 7c, e and 8). More specifically, MAD2L2 IRIF accumulation required the HEAT repeats of RIF1 and the same ATM/ATR phosphorylation sites in 53BP1

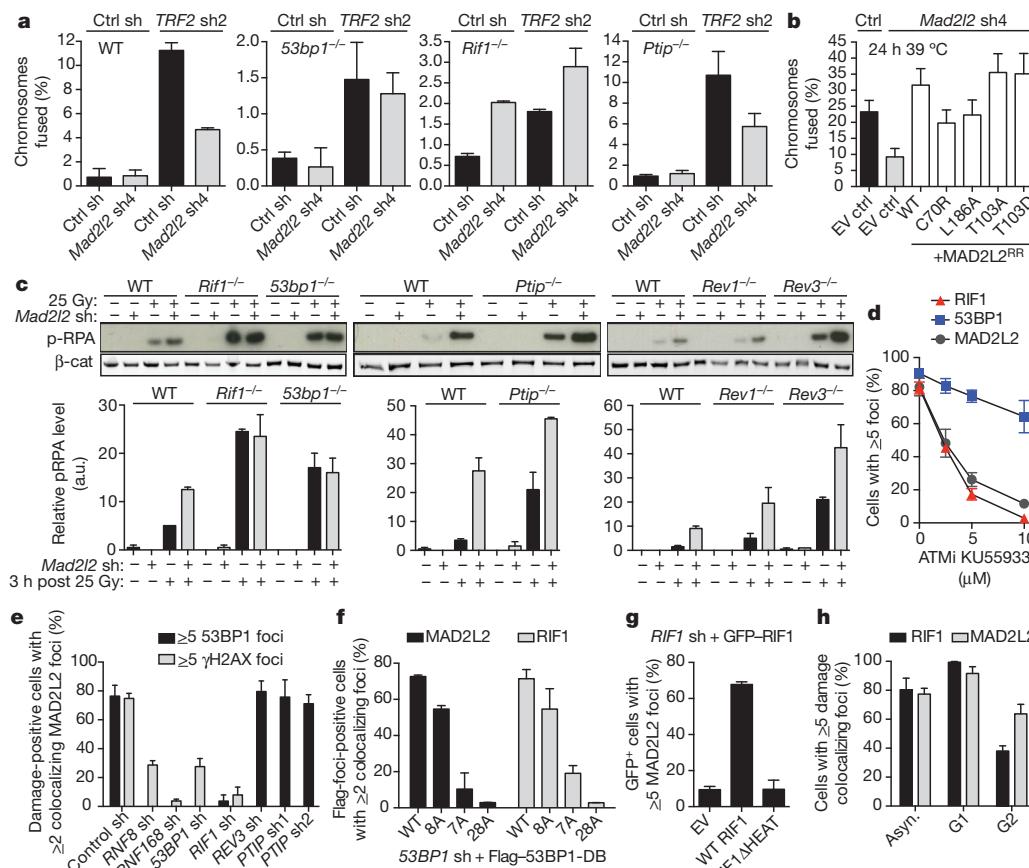


Figure 4 | MAD2L2 localizes to DSBs, inhibits end resection and promotes telomere-NHEJ in a 53BP1- and RIF1-dependent manner. **a**, Chromosome fusion after TRF2 inhibition in MEFs ($n = 2$, mean \pm s.e.m.). WT (wild-type) denotes controls to *53bp1*^{-/-} and *Rif1*^{-/-} MEFs. See Extended Data Fig. 6c for *Ptip*^{+/+} MEFs. **b**, Chromosome fusion in TRF2ts MEFs (mean \pm s.e.m.). EV, empty vector. **c**, p-RPA levels in MEFs ($n = 2$, mean \pm s.e.m.). a.u., arbitrary units. **d**, Irradiation-induced MAD2L2, RIF1 and 53BP1 foci in ATM inhibitor (ATMi)-treated U2OS cells ($n = 3$, mean \pm s.e.m.). **e**, MAD2L2 foci in irradiated shRNA-transduced U2OS cells ($n = 2$, mean \pm s.d.). **f**, MAD2L2 and RIF1 foci in irradiated U2OS cells expressing 53BP1-DB alleles with wild-type or substituted S/TQ sites (schematic in Extended Data Fig. 9a; $n = 2$, mean \pm s.d.). **g**, MAD2L2 foci in irradiated U2OS cells expressing wild-type RIF1 or RIF1 lacking heat repeats ($n = 2$, mean \pm s.d.). **h**, MAD2L2 and RIF1 foci in irradiated asynchronous, G1 or G2 synchronized RPE cells ($n = 3$, mean \pm s.d.).

that promote RIF1 IRIFs, while those that promote PTIP IRIFs were dispensable^{21–24} (Fig. 4f, g and Extended Data Fig. 9a–c). Similarly, MAD2L2 localization to uncapped telomeres required 53BP1 (Extended Data Fig. 8e). In line with dependency on RIF1, whose localization to DSBs is strongly reduced in G2, MAD2L2 foci were also diminished in G2, although less than RIF1 foci (Fig. 4h and Extended Data Fig. 9d). The latter could relate to a proposed role for MAD2L2 in the late steps of HDR²⁸.

In addition to inhibiting end resection, 53BP1, but not RIF1, promotes NHEJ by increasing mobility of uncapped telomeres^{13,20}. Consistent with MAD2L2 acting via end resection control downstream of RIF1, MAD2L2 deficiency did not impair mobility of uncapped telomeres (Extended Data Fig. 10a–c).

Altogether, our data identify MAD2L2 as an important regulator of DNA repair pathway choice that promotes NHEJ at telomeres and DSBs in multiple settings, by inhibiting CTIP- and EXO1-dependent 5' end resection downstream of RIF1 (Extended Data Fig. 10d). Whether MAD2L2 affects regulatory acetylation or phosphorylation of CTIP¹⁸, or blocks other end resection factors, awaits further investigation.

MAD2L2 and REV3 together compose Polζ and act with REV1 in translesion synthesis to facilitate DNA replication past blocking lesions⁴. On the basis of reduced gene conversion and impaired RPA foci, RAD51 foci and DSB resolution in REV1- and Polζ-depleted cells, it has been proposed that the ability of REV1/Polζ to alleviate replication stalling might also be important during HDR synthesis²⁸. Although requiring further characterization, together with the data presented here, this suggests differential roles for MAD2L2 in DSB repair. As shown here, MAD2L2 acts downstream of 53BP1/RIF1 in a REV1/REV3-independent manner to promote NHEJ by inhibiting resection, mostly in G1. In addition, MAD2L2, as part of Polζ, might promote HDR by facilitating replication of damaged DNA formed after resection in G2.

The contribution of MAD2L2 to DNA repair at uncapped telomeres and DSBs suggests that aberrant MAD2L2 expression could have pathological consequences by compromising genome integrity. In this respect it is interesting that in particular MAD2L2, more than REV1 or REV3, is reported to be overexpressed in several cancers and has prognostic significance (ref. 29 and the Oncomine database). Furthermore, our finding that MAD2L2 is essential for CSR identifies MAD2L2 as a potential disease-susceptibility gene for human primary immunodeficiencies.

Online Content Methods, along with any additional Extended Data display items and Source Data, are available in the online version of the paper; references unique to these sections appear only in the online paper.

Received 17 February; accepted 29 December 2014.

Published online 23 March 2015.

- Hanahan, D. & Weinberg, R. A. Hallmarks of cancer: the next generation. *Cell* **144**, 646–674 (2011).
- López-Otin, C., Blasco, M. A., Partridge, L., Serrano, M. & Kroemer, G. The hallmarks of aging. *Cell* **153**, 1194–1217 (2013).
- Zimmermann, M. & de Lange, T. 53BP1: pro choice in DNA repair. *Trends Cell Biol.* **24**, 108–117 (2014).
- Waters, L. S. et al. Eukaryotic translesion polymerases and their roles and regulation in DNA damage tolerance. *Microbiol. Mol. Biol. Rev.* **73**, 134–154 (2009).
- Konishi, A. & de Lange, T. Cell cycle control of telomere protection and NHEJ revealed by a ts mutation in the DNA-binding domain of TRF2. *Genes Dev.* **22**, 1221–1230 (2008).
- O'Sullivan, R. J. & Karlseder, J. Telomeres: protecting chromosomes against genome instability. *Nature Rev. Mol. Cell Biol.* **11**, 171–181 (2010).
- Davoli, T. & de Lange, T. The causes and consequences of polyploidy in normal development and cancer. *Annu. Rev. Cell Dev. Biol.* **27**, 585–610 (2011).
- Murnane, J. P. Telomere dysfunction and chromosome instability. *Mutat. Res.* **730**, 28–36 (2012).
- Smogorzewska, A., Karlsseder, J., Holtgreve-Grez, H., Jauch, A. & de Lange, T. DNA ligase IV-dependent NHEJ of deprotected mammalian telomeres in G1 and G2. *Curr. Biol.* **12**, 1635 (2002).

- Celli, G. B. & de Lange, T. DNA processing is not required for ATM-mediated telomere damage response after TRF2 deletion. *Nature Cell Biol.* **7**, 712–718 (2005).
- Peuschner, M. H. & Jacobs, J. J. DNA-damage response and repair activities at uncapped telomeres depend on RNF8. *Nature Cell Biol.* **13**, 1139–1145 (2011).
- Denchi, E. L. & de Lange, T. Protection of telomeres through independent control of ATM and ATR by TRF2 and POT1. *Nature* **448**, 1068–1071 (2007).
- Dimitrova, N., Chen, Y. C., Spector, D. L. & de Lange, T. 53BP1 promotes non-homologous end joining of telomeres by increasing chromatin mobility. *Nature* **456**, 524–528 (2008).
- Deng, Y., Guo, X., Ferguson, D. O. & Chang, S. Multiple roles for MRE11 at uncapped telomeres. *Nature* **460**, 914–918 (2009).
- Dimitrova, N. & de Lange, T. Cell cycle-dependent role of MRN at dysfunctional telomeres: ATM signaling-dependent induction of nonhomologous end joining (NHEJ) in G1 and resection-mediated inhibition of NHEJ in G2. *Mol. Cell. Biol.* **29**, 5552–5563 (2009).
- Orthwein, A. et al. Mitosis inhibits DNA double-strand break repair to guard against telomere fusions. *Science* **344**, 189–193 (2014).
- Daniel, J. A. & Nussenzweig, A. The AID-induced DNA damage response in chromatin. *Mol. Cell* **50**, 309–321 (2013).
- Panier, S. & Durocher, D. Push back to respond better: regulatory inhibition of the DNA double-strand break response. *Nature Rev. Mol. Cell Biol.* **14**, 661–672 (2013).
- Sfeir, A. & de Lange, T. Removal of shelterin reveals the telomere end-protection problem. *Science* **336**, 593–597 (2012).
- Zimmermann, M., Lottersberger, F., Buonomo, S. B., Sfeir, A. & de Lange, T. 53BP1 regulates DSB repair using Rif1 to control 5' end resection. *Science* **339**, 700–704 (2013).
- Di Virgilio, M. et al. Rif1 prevents resection of DNA breaks and promotes immunoglobulin class switching. *Science* **339**, 711–715 (2013).
- Chapman, J. R. et al. Rif1 is essential for 53BP1-dependent nonhomologous end joining and suppression of DNA double-strand break resection. *Mol. Cell* **49**, 858–871 (2013).
- Escribano-Díaz, C. et al. A cell cycle-dependent regulatory circuit composed of 53BP1-Rif1 and BRCA1-CTIP controls DNA repair pathway choice. *Mol. Cell* **49**, 872–883 (2013).
- Callen, E. et al. 53BP1 mediates productive and mutagenic DNA repair through distinct phosphoprotein interactions. *Cell* **153**, 1266–1280 (2013).
- Hara, K. et al. Crystal structure of human REV3 in complex with a human REV3 fragment and structural implication of the interaction between DNA polymerase zeta and REV1. *J. Biol. Chem.* **285**, 12299–12307 (2010).
- Khalaj, M. et al. A missense mutation in Rev7 disrupts formation of Polzeta, impairing mouse development and repair of genotoxic agent-induced DNA lesions. *J. Biol. Chem.* **289**, 3811–3824 (2014).
- Chan, K. L., Palmai-Pallag, T., Ying, S. & Hickson, I. D. Replication stress induces sister-chromatid bridging at fragile site loci in mitosis. *Nature Cell Biol.* **11**, 753–760 (2009).
- Sharma, S. et al. REV1 and polymerase zeta facilitate homologous recombination repair. *Nucleic Acids Res.* **40**, 682–691 (2012).
- Rimkus, C. et al. Expression of the mitotic checkpoint gene MAD2L2 has prognostic significance in colon cancer. *Int. J. Cancer* **120**, 207–211 (2007).

Acknowledgements We thank A. Konishi, G. Celli and T. de Lange for TRF2(Ile468Ala), *Trf2*^{fl/fl}–*p53*^{+/−} MEFs and *Trf2*^{fl/fl}–*p53*^{+/−}*Lig4*^{−/−} MEFs, P. Bouwman, M. Pieterse, T. Halazonetis, O. Kallioniemi, B. Gerritsen, B. Morris and P. Halonen for composing and providing the DDR TRC shRNA sub-library, members of the NKI Genomics Core Facility for technical support, A. van Kessel for *MAD2L2* cDNA, E. Hendrickson for repair plasmids, R. Chapman and S. Boulton for mRif1 antibody, *53bp1*^{−/−} and *Rif1*^{−/−} MEFs, E. Callen and A. Nussenzweig for 53BP1 and PTIP expression vectors and for *Ptip*^{−/−} MEFs, N. de Wind for *Rev1*^{−/−} and *Rev3*^{−/−} MEFs, B. van den Brook for help with telomere mobility analysis, I. de Krijger, Z. Yalcin and M. Simonetta from the Jacobs group for contributing to end resection and interaction analysis and members of the R. Medema laboratory for discussion. This work was supported by grants from the European Research Council (ERCStG 311565) and Dutch Cancer Society (KWF-NKI2007-3907) to J.J.L.J., a grant from the Canadian Institutes for Health Research (CIHR MOP89754) to D.D., a La Caixa fellowship to S.S.B. and a CIHR post-doctoral fellowship to A.O. D.D. is a Canada Research Chair (tier 1) in the Molecular Mechanisms of Genome Integrity and J.J.L.J. is an EMBO Young Investigator.

Author Contributions J.J.L.J., V.B., N.M. and S.S.-B. designed the experiments and analysed the data. M.H.P. and J.v.d.T. performed the TIGR-screen, analysed the screen results and performed initial validation. V.B., N.M., S.S.-B. and B.A.W. performed all experiments, with the exception of assessment of CSR and sister-telomere fusion upon repair activation in mitosis, which were performed and analysed by A.O. and D.D. S.S.-B., M.H.P. and J.v.d.T. contributed equally. J.J.L.J. wrote the manuscript.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to J.J.L.J. (j.jacobs@nki.nl).

METHODS

TIGIR screen. *Trf2^{-/-}*–*p53^{-/-}* TRF2ts MEFs (in short TRF2ts MEFs) were infected in triplicate with PLKO lentiviral shRNA pools consisting of 1,976 shRNAs selected from the Mouse TRC 1.0 Collection. Cells were infected with a multiplicity of infection below 1 and drug selected on 4 µg ml⁻¹ puromycin. After completion of selection, TRF2ts MEFs were plated in duplicate on 15-cm dishes at 150,000 cells per dish and allowed to adhere at 32 °C. One set of plates was kept at 32 °C for 4 additional days, the other set of plates was subjected to 12 days of telomere uncapping at 39 °C, followed by recovery for 4 additional days at 32 °C. Cells were collected for genomic DNA isolation using a DNeasy Blood & Tissue Kit (Qiagen) and shRNA inserts were recovered from genomic DNA by 2 sequential PCR amplifications (PCR1 and PCR2). In PCR1 6 individual PCR reactions were performed on 325 ng genomic DNA using for each sample a common reverse primer P7_PLKO1_R (5'-CAAGCAGAAGACGGCATACGAGATTCTTCCCTGACTGTACCC-3') and 1 of 6 forward primers with a unique index (underlined): Illuseq_01_PLKO1_f (5'-ACACTTTCCCTACACGACGCTTCCGATCTCGTACITGTGG AAAGGACGAAACACCGG-3'), Illuseq_02_PLKO1_f (5'-ACACTTTCCCT ACACGACGCTTCCGATCTACATCGCTGTGGAAAGGACGAAACACC GG-3'), Illuseq_03_PLKO1_f (5'-ACACTTTCCCTACACGACGCTTCCG ATCTGCCTAACATTGTGGAAAGGACGAAACACCGG-3'), Illuseq_04_PLKO1_f (5'-ACACTTTCCCTACACGACGCTTCCGATCTGTGGCACTTG TGAAAGGACGAAACACCGG-3'), Illuseq_05_PLKO1_f (5'-ACACTTTCCCT ACACGACGCTTCCGATCTACTGTCTGTGGAAAGGACGAAACACC GG-3'), Illuseq_06_PLKO1_f (5'-ACACTTTCCCTACACGACGCTTCCG ATCTATTGGCTTGTGGAAAGGACGAAACACCGG-3'). In PCR2, 2.5 µl DNA from each PCR1 reaction was amplified using primers P5_Illuseq (5'-AAT GATA CGGC GACC ACCG GAG ATCT A CACT TTCC CT ACAC GAC GCT TCC CCG AT CT-3') and P7 (5'-CAAGCAGAAGACGGCATACGAGAT-3'). PCR products were pooled per condition, purified using a Qiagen PCR purification kit (Qiagen), analysed on a bioanalyzer for quality and concentration, and subjected to deep-sequencing on a Illumina Hiseq2000 genome analyser at the NKI Genomics Core Facility.

Cell culture, growth assays, flow-cytometry. *Trf2^{-/-}*–*p53^{-/-}* TRF2ts MEFs (TRF2ts MEFs) and *Trf2^{-/-}*–*p53^{-/-}*–*Lig4^{-/-}* TRF2ts MEFs (TRF2ts *Lig4^{-/-}* MEFs) were generated from *Trf2^{fl/fl}*–*p53^{-/-}* MEFs and *Trf2^{fl/fl}*–*p53^{-/-}*–*Lig4^{-/-}* MEFs as described before¹¹. All cells were grown in DMEM with 100 U penicillin, 0.1 mg ml⁻¹ streptomycin, 2 mM L-glutamine and 10% FBS. TRF2ts MEFs were maintained at the permissive temperature of 32 °C and only grown at 37 °C or 39 °C to induce telomere uncapping through inactivation of TRF2. All other cells were grown at 37 °C. Wild-type, 53bp1^{-/-}, *Rif1*^{-/-}, *Ptip*^{+/+} and *Ptip*^{-/-} MEFs were SV40 immortalized.

For colony survival assays, TRF2ts MEFs were plated at 150,000 cells per 15-cm dish or 45,000 per 10-cm dish, allowed to adhere and kept at 32 °C or placed at 39 °C. To evaluate potential toxicity or growth rate differences, plates were fixed with 10% formalin and stained with 0.1% crystal violet after growth for 1 week at 32 °C. To evaluate survival upon prolonged telomere uncapping plates were fixed and stained after growth for 12 days at 39 °C, followed by 0 or 2 weeks recovery at 32 °C. Quantifications of survival assays reflect the relative survival after growth for 12 days at 39 °C and 2 weeks recovery at 32 °C, corrected for plating efficiency and growth at 32 °C and assessed by crystal violet extraction as described below.

To address cell survival under telomere uncapping conditions in short-term growth assays, TRF2ts MEFs were plated in quadruplicate at 5,000 cells per well on a 12-well plate (Fig. 1c) or in duplicate at 2,500 cells per well on a 24-wells plate (Extended Data Fig. 5h), allowed to adhere overnight at 32 °C and then placed for 12–14 days at 39 °C. During this time period of non-permissive temperature, plates were fixed every 2–4 days with 10% formalin and stained with 0.1% crystal violet. Crystal violet was extracted from the plates using 10% acetic acid and its absorbance at 595 nm was quantified using a Tecan microplate reader. Quantifications were corrected for plating efficiency.

Cell cycle distribution of TRF2ts MEFs with or without MAD2L2 knockdown was determined by propidium iodide staining, as well as by BrdU incorporation, and acquired on a BD Fortessa using FACSDiva software (BD Biosciences) and analysed with FlowJo (TreeStar) software.

Analysis of tetraploidy (8N) and aneuploidy (>4N) was based on DNA-content analysis by flow cytometry of TRF2ts cells grown at 32 °C or 39 °C for 48 h and then stained with propidium iodide.

Cell cycle synchronisation was done as follows. For G1 cells, RPE cells were kept in 2.5 mM thymidine (T1895, Sigma) for 16 h, released for 6 h, treated with 10 µM RO3306 (Calbiochem) for 16 h. Synchronised cells were released for 1 h, mitotic cells were collected by mitotic shake off and replated to let them continue into G1 3.5 h before irradiation.

For G2 cells, RPE cells were kept in 2.5 mM thymidine for 16 h, released for 6 h and treated with 2.5 mM thymidine for 16 h. Synchronised cells were released for 4.5 h before irradiation to let cells progress to G2.

Plasmids. Cells were transduced as before¹¹ with the following pLKO-puro shRNA lentiviruses obtained from Mission library clones (Sigma) against mouse genes: *Mad2l2* sh1 (TRCN0000012843: 5'-CCAGACTCCAAGTGCCTTAT-3'), *Mad2l2* sh2 (TRCN0000012844: 5'-CCAGTGGAGAAGTTGTCTTT-3'), *Mad2l2* sh3 (TRCN0000012845: 5'-CGAGCCTTCATCCTTAAGATT-3'), *Mad2l2* sh4 (TRCN0000012846: 5'-CATCTTCCAGAAGCGCAAGAA-3'), *Rev1* sh pool (pool of TRCN00000120297: 5'-CAGCAGTGCCTGTGAGGTATT-3', TRCN0000120298: 5'-GCCGAGATCAACTATGGAATA-3', TRCN0000120299: 5'-GCTGGAATG AAGACGGTGTAA-3', TRCN0000120300: 5'-GCCCTCCGATTGAAATCA AA-3', TRCN0000120301: 5'-GCACGTTGATATGGACTGCTT-3'), *Rev1* sh1 (TRCN00000120298: 5'-GCCGAGATCAACTATGGAATA-3'), 53bp1 sh1 (TRCN0000081778: 5'-GCTATTGTGGAGATTGTGTT-3' and against human genes: *MAD2L2* sh (TRCN0000006570: 5'-CCCGGAGCTGAATCAGTATAT-3'), 53BP1 sh (TRCN0000018865: 5'-GATACTTGGCTTACTGGTT-3'), *RIF1* sh (TRCN0000155431: 5'-CGCATTCTGCTGTGTTGATT-3'), *PTIP* sh1 (TRCN0000129174: 5'-CTTCAGATTATCACCAGAA-3'), *PTIP* sh2 (TRCN0000322819: 5'-GCCATGTTCACAGCATATTAT-3'), *RNF8* sh: (TRCN0000003439: 5'-GA AGCCGTTATGAATGTGAAA-3'), *RNF168* sh (TRCN0000034137: 5'-GCA GTAGTTAATAGAAGAAA-3'), *REV3* sh (TRCN0000244436: 5'-TGACCTGTC TGAGACTATTTA-3'), *RAD51* sh (TRCN0000018879: 5'-CGGTCAAGAGATC ATACAGATT-3').

For knockdown of *Ctip* and *Exo1* TRF2ts MEFs were transduced as before¹¹ with pRetroSuper-blas shRNA retroviruses containing the following target sequences: *Ctip* sh1: 5'-GCAAGGTTACAAGTCAAAGT-3', *Ctip* sh2: 5'-GCAGACCTT TCTCACTATA-3', *Ctip* sh3: 5'-GCATTAACCGGCTACGAAA-3', *Exo1* sh1: 5'-GCATTGGCACAAGAATTAA-3'. Retroviral shRNAs used for knockdown of *Rnf8*, *Lig4* or *TRF2* have been described before¹¹. The 53bp1 shRNA vector used in CSR assays was obtained from Origene (FI334760) and contained the following target sequence: 5'-GGAAGTTGAAACCAAGTGTGATTAGTATTG-3'.

For localization and complementation experiments human *MAD2L2* complementary DNA was Gateway-cloned (Invitrogen) into pMSCV-retroviral vectors containing coding sequences for amino-terminal Flag or GFP epitopes. *MAD2L2* deletion constructs were generated through Gateway-cloning (Invitrogen). The target sequences of *Mad2l2* sh4 or the human *MAD2L2* shRNA were changed to 5'-GATTTCAGAACGCAAGAA-3' or 5'-TCCCGAACTGAATCAGTATA T-3', respectively, to create RNAi-resistant epitope-tagged *MAD2L2*^{RR} in pMSCV. GFP-*RIF1*, Flag-53BP1-DB and PTIP-GFP expression vectors were described before^{23,24}. The target sequence of the human *RIF1* shRNA was changed to 5'-GCA GCGAAGTTGAAACTTGA-3' to create RNAi-resistant GFP-*RIF1*.

Quantitative real-time PCR. RNA was isolated using TRIzol reagent (Ambion) and reverse transcribed into cDNA using AMV first-strand cDNA synthesis kit for RT-PCR (Roche) according to the manufacturer's instructions. Quantitative real-time PCR (qRT-PCR) on cDNA was performed in triplicate using Power SYBR Green PCR Master Mix (Applied Biosystems) on the StepOnePlus real-time PCR system.

The following primers were used: for mouse *Mad2l2* forward 5'-ACACTCCAC TCGCTCAACCC-3' and reverse 5'-AAAGACAACTTCTCCACTGGGC-3'; for human *MAD2L2* forward 5'-CGAGTTCCTGGAGGTGGCTGTGCATC-3' and reverse 5'-CTTGACCCAGTCAGCCTGTCCTGGATA-3'; for human *RNF8* forward 5'-TGAAGAGCTAACATCGCAGCA-3' and reverse 5'-TCAAGGTG ACAGCCTCAATGAA-3'; for human *RNF168* forward 5'-GGCGAGTTATGC TGTCCT-3' and reverse 5'-GCCGCCACCTTGCTTATTTC-3'; for mouse *Rev1* forward 5'-AAGCCAGATGGTCAGTACCAAC-3' and reverse 5'-TCCATTGA ACGTCCAATCCCG-3'; for human *REV3* forward 5'-AGGACTCGAAGTCA CCTATGC-3' and reverse 5'-AGAGGTAACCCAGGAATGC-3'; for mouse *Ctip* forward 5'-GTGCTGGTCCGGAGCAG-3' and reverse 5'-TTGACTTGTAAAC CTTGCACTTCC-3'; for mouse *Exo1* forward 5'-TAAACACGTCGAGCCTGT CC-3' and reverse 5'-CAGAGCCCAGGAACCTTGT-3'; for human *RAD51* forward 5'-GTGGAGCTAATGGCAATGCAG-3' and reverse 5'-TTAGCTCT TCTTTGGCGCA-3'; for mouse *Hprt* forward 5'-CTGGTAAAGGACCTCTC G-3' and reverse 5'-TGAAGTACTCATTATAGTCAGGAGTC-3'; for human *GAPDH* forward 5'-GAAGGTGAAGGTCGGAGTC-3' and reverse 5'-GAAGA TGGTGATGGGATTTC-3'.

Immunoblotting. Whole-cell lysates were prepared by scraping cells in SDS sample buffer; protein concentration was determined by standard BCA protein assay (Pierce) and equal amounts of protein were separated on precast 4–12% Bis-Tris or 3–8% Tris-acetate gels (Invitrogen). Immunoblotting was done according to standard methods using IRDye800CW- and IRDye680-labelled secondary antibodies for detection on the Odyssey Infrared imager (LI-COR) or using horseradish peroxidase (HRP)-conjugated secondary antibodies for detection by enhanced chemiluminescence (Supersignal, Thermo Scientific). Primary antibodies used were against *MAD2L2* (sc135977, Santa Cruz, 1:500), phospho-ATM S1981 (4526, Cell Signaling, 1:1,000), phospho-KAP1 S824 (Bethyl, 1:1,000), 53BP1 (NB100-305, Novus, 1:1,000).

1:500), phospho-H2AX S139 (clone JBW301, Upstate, 1:1,000), CHK2 (611570, BD, 1:1,000), phospho-CHK1 S345 (2348, Cell Signaling, 1:1,000), phospho-ATR S428 (Cell Signaling, 1:1,000); RPA34 (RPA34-20, Gene Tex, 1:1,000); phospho-RPA32 S4/S8 (A300-245A, Bethyl, 1:1,000); CTIP (sc22838, Santa Cruz, 1:500); γ -tubulin (GTU-88, Sigma, 1:5,000); β -catenin (610154, BD Biosciences, 1:2,000); GFP (a-11122, Invitrogen, 1:500); RIF1 (sc55979, Santa Cruz, 1:1,000); PTIP (ABE69, Millipore and ab2614, Abcam, 1:500); actin (CP01, Calbiochem, 1:5,000). Immunoblotting for tubulin, β -catenin or actin served as loading controls.

Immunofluorescence. For immunofluorescence, cells were grown on 8-well chamber slides (Lab-Tek) and in cases of MAD2L2 detection first pre-extracted for 30 s to 5 min with 0.5% Triton/PBS on ice, then fixed for 10 min in 4% paraformaldehyde in PBS. Cells were permeabilized for 5 min (MEFs) or 10 min (U2OS) in 0.5% Triton/PBS, washed twice in PBS, incubated for 1 h in blocking solution (0.02% Triton, 5% NGS, 5% FCS in PBS) and overnight at 4 °C with primary antibodies in blocking solution. Immunofluorescence for PTIP-GFP in U2OS was done as described²⁴. Primary antibodies used were against phospho-ATM S1981 (4526, Cell Signaling, 1:1,000), phospho-H2AX Ser 139 (05-636, Millipore and 2577S, Cell Signaling, 1:500), 53BP1 (A300-272A, Bethyl, 1:2,000), mRIF1 (gift from R. Chapman, 1:500)²², RIF1 (A300-569A, Bethyl, 1:1,000), MAD2L2 (sc135977, Santa Cruz, 1:100), Flag (F3165 or F1804 at 1:500, F7425 at 1:400, Sigma), GFP (a-11122, Invitrogen, 1:500). Cells were washed three times with 0.02% Triton/PBS, followed by incubation with Alexa Fluor 488 or 568 goat anti-mouse or anti-rabbit IgG secondary antibodies (Invitrogen) in blocking solution for 1 h. After four washes in 0.02% Triton/PBS, slides were mounted in Vectashield (Vector Laboratories) containing DAPI. Confocal fluorescence images were obtained on a Leica SP5 confocal system equipped with an Ar, Kr and HeNe laser system. Images were taken with a \times 63 NA 1.32 oil objective and standard LAS-AF software. Possible crosstalk between fluorochromes was avoided by careful selection of imaging conditions. A minimum of 100 cells per condition per experiment was analysed. DDR foci in MEFs were captured using the Metafer4/MetaCyte platform (MetaSystems) equipped with an AxioImager Z2 microscope (Carl Zeiss). Images of random selections of cells were acquired with an EC ‘Plan-Neofluar’ \times 40/0.75 objective. Analysis was done using MetaCyte software, carefully fine-tuned for each antibody. A minimum of 500 cells was analysed per condition per experiment for p-ATM, γ H2AX, 53BP1 and RIF1. All assessments of irradiation-induced DDR foci were done 3 h after irradiation with 5 Gy, with the exception of Fig. 3a, which is 30 min after 10 Gy.

Metaphase chromosome analysis. Cell collection, preparation of metaphase spreads and telomere FISH with a FITC-OO-(CCCTAA)₃ peptide nucleic acid custom probe (Biosynthesis) for metaphase chromosome analysis was done as described⁵.

Digital images of metaphases were captured using the Metafer4/MSearch automated metaphase finder system (MetaSystems) equipped with an AxioImager Z2 microscope (Carl Zeiss). After scanning metaphase preparations at \times 10 magnification, high-resolution images of metaphases were acquired using a ‘Plan-Apochromat’ \times 63/1.40 oil objective. Chromosome fusions were quantified from >30 metaphases.

For monitoring mitotic telomeres by PNA FISH, IMR90 (E6/E7) cells were transduced with pMX-Flag-53BP1 (1-1711) or its mutant Thr1608Ala/Ser1618Ala (TASA) and pMX-GFP-RNF8 or its mutant Thr198Ala and selected with 0.2 μ g ml⁻¹ puromycin, for 3 days before plating. Cells were transfected with siRNAs targeting endogenous RNF8 (ThermoFisher, D-006900-01, target sequence: 5'-AGAAUG AGCUCCAUGUAUUU-3') and 53BP1 (ThermoFisher, D-003548-01, target sequence: 5'-GAGAGCAGAUGAUCCUUUA-3') and the following siRNAs: PTIP siRNA pool (ThermoFisher, D-012795-01-0010, target sequences: 5'-UUAAAG CGAGCAAGUAUUA-3', 5'-GAGCCUGGGUUGAUUAACU-3', 5'-GUUGA UGAGUUAAGACUA-3' and 59-GCCAUCAGCUGCUGUUUA-3'), RIF1 (ThermoFisher, D-027983-02, target sequence: 5'-AGACGGUCUCUUAUGU UA-3'), MAD2L2 (ThermoFisher, D-003272-10: 5'-GAGAAUUCGUCUUUG AGA-3') and non-targeting control siRNA (ThermoFisher, D-001210-02, target sequence: 5'-UAAGGCUAUGAAGAGAUAC-3') using RNAiMax (Invitrogen) and synchronized with a double-thymidine block. Cells were subsequently processed as previously described¹⁶.

Pulsed-field gel electrophoresis and in-gel detection of telomeric DNA. Analysis of 3' single-stranded G-overhangs at mouse telomeres was performed as described⁵ by pulsed-field gel electrophoresis and in-gel hybridization of a ³²P-labelled telomeric repeat (CCCTAA)₄ oligonucleotide to native DNA. In brief, $Trf2^{-/-}$ p53^{-/-} TRF2ts MEFs were collected at the permissive temperature as well as after growth at the non-permissive temperature. Agarose plugs were prepared containing 1×10^6 cells per plug and digested overnight with 1 mg ml⁻¹ proteinase K at 50 °C. Before loading on a 1% agarose gel, plugs were digested overnight with 60 U of MboI enzyme per plug. After running the gel on a CHEF-DR III pulsed-field gel electrophoresis system (BioRad), the gel was dried, pre-hybridized in Church Mix containing 7% SDS and hybridized overnight with a ³²P-labelled (CCCTAA)₄ probe, all under non-denaturing conditions. Signal was captured on a phosphorimager

and analysed with AIDA Image Analyzer software. To confirm that the G-overhang signal detected with this method was indeed derived from 3' single-stranded telomeric TTAGGG repeats, digestion with 3' exonuclease and hybridization with a ³²P-labelled (TTAGGG)₄ probe specific for the C-rich strand was performed on separate plugs. After capturing the hybridization signal from single-stranded telomeric DNA, the gel was denatured and re-hybridized with the same probes to obtain a total telomere signal for the purpose of quantification.

Repair and clonogenic survival assays. To quantify the repair of DSBs by NHEJ, U2OS cells transduced with shRNAs were transfected with pEJ5-GFP (NHEJ), as well as an mCherry expression vector and pCAGGS-Isce-I-Trex2 (ref. 30). Transfections were done using Mirus TransIT LT-1 (Mirus Bio LLC). The fraction of mCherry-positive cells that was GFP-positive was determined 72 h after transfection by flow-cytometry on a BD LSRIFortessa using FACSDiva software. Quantification was carried out with FlowJo software (TreeStar).

For random plasmid integration assays 500,000 U2OS cells were seeded in 6-cm dishes 24 h before transfection with 3.9 μ g of NdeI-linearized pMSCVblas-GFP and 0.1 μ g of mCherry plasmids using ViaFect (Promega). Transfected cells were trypsinized the following day and seeded in 10-cm dishes at different densities for colony formation. Selection was initiated the following day with blasticidin at 5 μ g ml⁻¹. Cells were fixed with 4% paraformaldehyde 10–14 days after plating and stained with 0.1% crystal violet. Colony counting was performed on a Col-Count (Oxford Optronix) and values were normalized for plating and transfection efficiencies.

For clonogenic survival assays U2OS cells were seeded in duplicate in 6-well plates at 100, 200 and 400 cells per well for unirradiated cells, and at 400, 800 and 1,600 cells per well for irradiated cells. Cells were grown for <20 h before γ -irradiation and then fixed after 10–12 days of recovery with 4% paraformaldehyde and stained with 0.1% crystal violet. Colonies with >50 cells were counted on a GelCount (Oxford Optronix) and the surviving fraction was determined.

CSR assays. Research involving animals was performed in accordance with protocols approved by the animal facility at Toronto Centre for Phenogenomics (Toronto). Primary B cell extraction was performed as described before²³. In brief, mature B lymphocytes were isolated from spleens of 8–12-week-old C57BL/6 mice by depletion of CD43⁺ cells using CD43 microbeads (Miltenyi Biotech) according to the manufacturer's instructions. Purified B cells were resuspended at a concentration of 10^6 cells per ml in the presence of 25 μ g ml⁻¹ LPS (Sigma-Aldrich) to allow B cell proliferation. Both mouse CH12F3-2 (ref. 31) and primary B cells were infected with lentiviral/retroviral particles obtained from the supernatant of HEK293T cell cultures cotransfected with the lentiviral/retroviral shRNA constructs and the vector expressing MLV Gag-Pol (for retrovirus) or the lentiviral psPAX2 packaging vector with VSV-G envelope (2:1 ratio). In brief, the viral supernatant was concentrated at 25,000g for 90 min at 25 °C before resuspension in RPMI culture medium. Infection was subsequently performed as described²³ and followed by selection with puromycin.

IgM to IgA switching was assayed in CH12F3-2 cells activated with 1 ng ml⁻¹ TGF- β 1 (Preprotech), 10 ng ml⁻¹ IL-4 (Preprotech) and 1 μ g ml⁻¹ agonist anti-CD40 (BD). IgA expression was measured by flow-cytometry using biotin-conjugated anti-mouse IgA antibody (1:200; eBioscience) followed by an anti-biotin-APC antibody at 0, 24 and 48 h after activation. Proliferation of the different transduced CH12F3-2 cell lines was monitored using CFSE (Invitrogen) following the manufacturer's guidelines. CSR assays in both cell types were done in triplicate for every independent experiment.

Ex-vivo isotype switching to IgG1 in primary B cells was assayed by flow-cytometry using an anti-IgG1-biotin (1:200; BD) followed by allophycocyanin (APC)-conjugated anti-biotin (Miltenyi Biotech) and 10 μ g ml⁻¹ propidium iodide 5 days after isolation.

For semiquantitative RT-PCR analysis of *Aid* mRNA, IgM and IgA germ-line transcripts total RNA was extracted with TRIzol (Invitrogen) and cDNA was synthesized by random priming using the Superscript III reverse transcriptase kit (Invitrogen). Mouse *Aida* transcripts were amplified using oligonucleotides 5'-G TGCCACCTCCTGCTCACTGG-3' and 5'-TTCATGTAGCCCTCCCAGGC-3'. IgM (μ) germ-line transcripts were amplified using oligonucleotides 5'-AAAATG TCCGCTGGCTAAC-3' and 5'-AGAACAGTCAGTGTAGGCAGTAGA-3', IgA (I α) using 5'-CCTGGCTGTTCCCTATGAA-3' and 5'-GAGCTGGTGGG AGTGTCACTG-3'; and *Gapdh* transcripts were amplified using oligonucleotides 5'-AGCCTCGCCGTAGACAA-3' and 5'-AATCTCCACTTGCCACTGC-3'. 53BP1 depletion is known to impair CSR³ and was included as control.

TSCE. TSCE were detected by chromosome-orientation FISH, essentially as described³² with minor modifications. In brief, for chromosome-orientation FISH cells were grown in 10 μ M BrdU:BrdC (3:1) for 24 h with the addition of 0.2 μ g ml⁻¹ demeclocine for the final 2 h. Slides were treated with RNase A (0.5 mg ml⁻¹) for 10 min at 37 °C, stained with Hoechst 33258 (0.5 μ g ml⁻¹) in 2 \times SSC for 15 min at room temperature and exposed to 5.4×10^3 J m⁻² 365-nm ultraviolet light

(Stratalinker 2400 UV irradiator). Following digestion with exonuclease III ($10\text{ U}\mu\text{l}^{-1}$, Promega) for 10 min at room temperature, slides were dehydrated through an ethanol series (70%, 95% and 100%) and incubated sequentially with Cy3-TelG 5'-[TTAGGG]₃-3' and FITC-TelC 5'-[CCCTAA]₃-3' probes at room temperature.

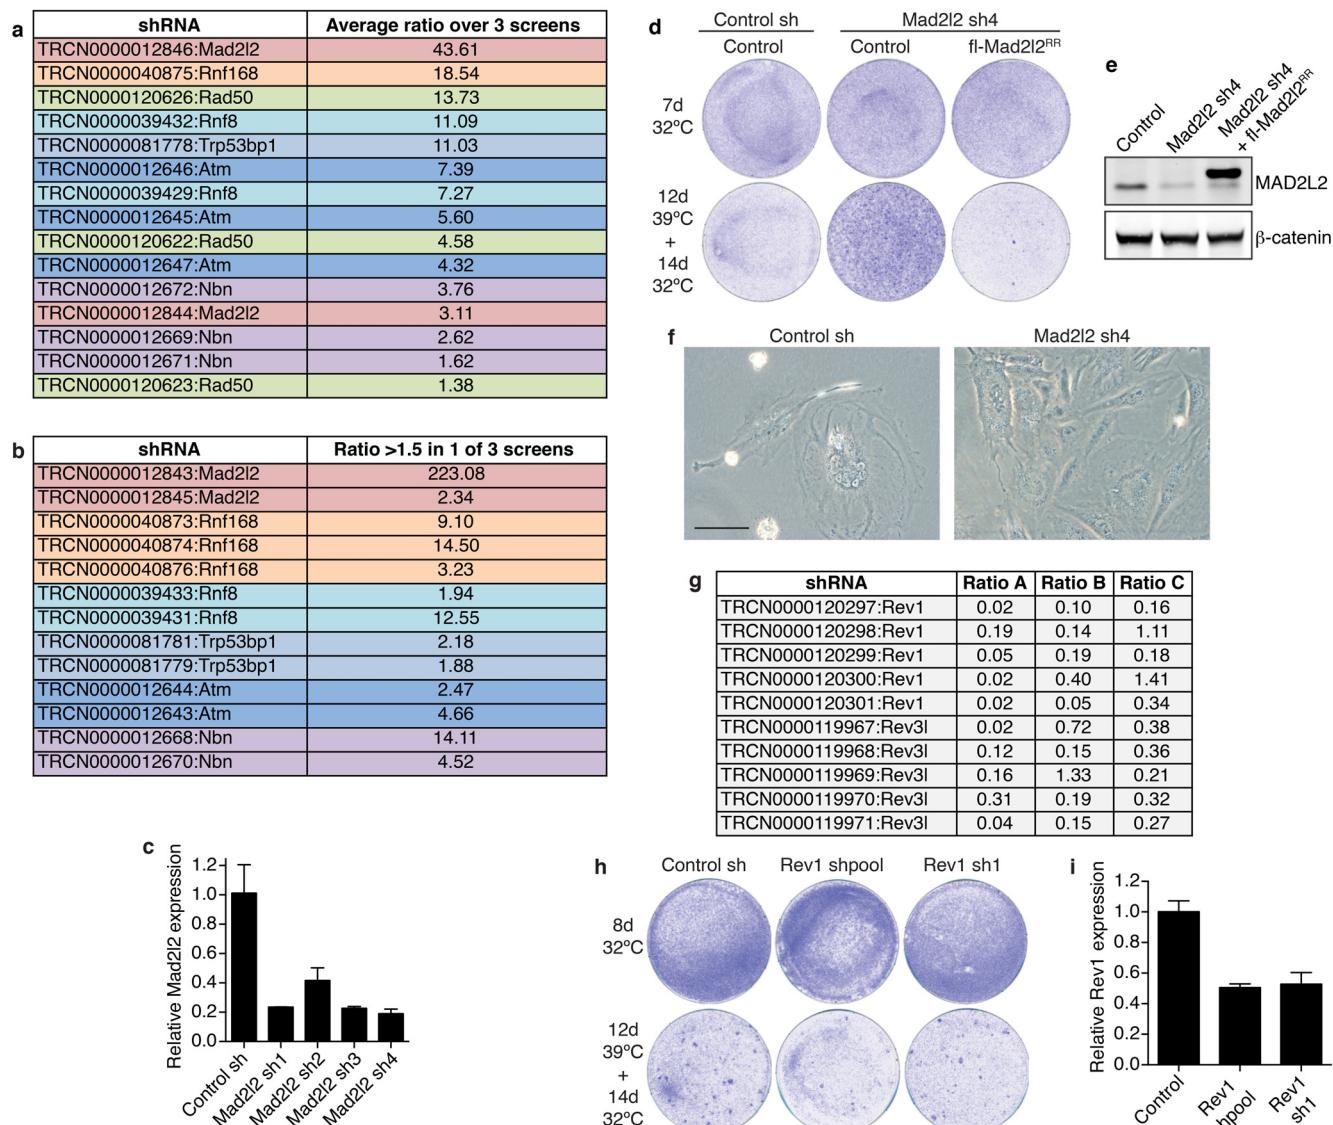
Telomere mobility. Telomere mobility analysis was done as described¹³ with minor adjustments. TRF2ts cells expressing eGFP-TRF1 and mCherry-BP1-2 (53BP1 amino acids 1220–1711) were seeded on Willco wells and left to attach overnight at 32 °C. Before imaging, cells were incubated at 37 °C for 30 min in Leibovitz's L15 medium inside the Microscope climate chamber. Cells were monitored every minute for 20 min ($t = 20$ frames) at 37 °C. Five-micrometre Z-stacks at 0.5-μm steps in both eGFP and mCherry channels were obtained with a $\times 63$ 1.4 oil objective using ZEN software on an Observer Z1 microscope system (Carl Zeiss). Images were obtained at 2×2 binning with 512×512 pixels. Tracking of telomeres was performed with FIJI software, cells were registered with the Rigid body option from the StackReg plugin. Particles were then tracked with the particle detector and

tracker plugin from Mosaic. Telomeres to track were selected on basis of colocalization with mCherry-BP1-2, and tracked for at least 17 out of 20 frames.

Database. To address whether changes in *MAD2L2* mRNA expression occur in human cancers we consulted the Oncomine database at <https://www.oncomine.org/resource/login.html>.

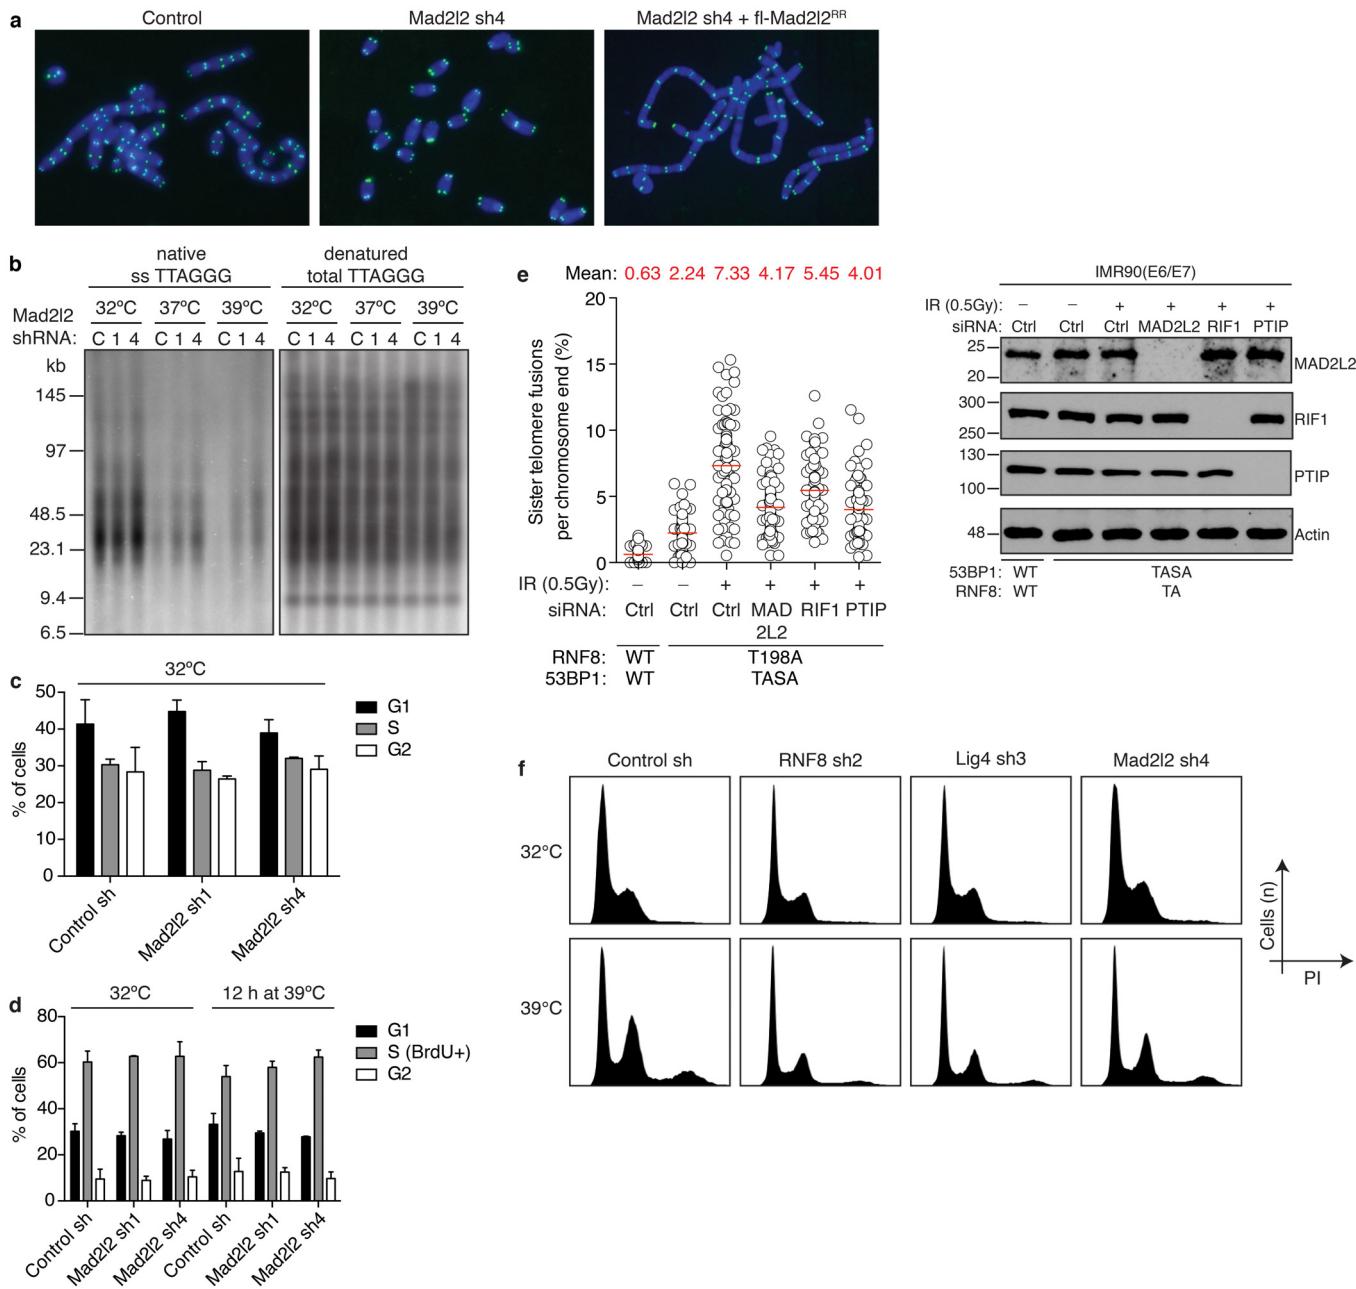
Statistics. No statistical methods were used to predetermine sample size. *P* values were determined by Kruskal–Wallis ANOVA test.

30. Bennardo, N., Gunn, A., Cheng, A., Hasty, P. & Stark, J. M. Limiting the persistence of a chromosome break diminishes its mutagenic potential. *PLoS Genet.* **5**, e1000683 (2009).
31. Kato, H. et al. Involvement of RBP-J in biological functions of mouse Notch1 and its derivatives. *Development* **124**, 4133–4141 (1997).
32. Celli, G. B., Denchi, E. L. & de Lange, T. Ku70 stimulates fusion of dysfunctional telomeres yet protects chromosome ends from homologous recombination. *Nature Cell Biol.* **8**, 885–890 (2006).



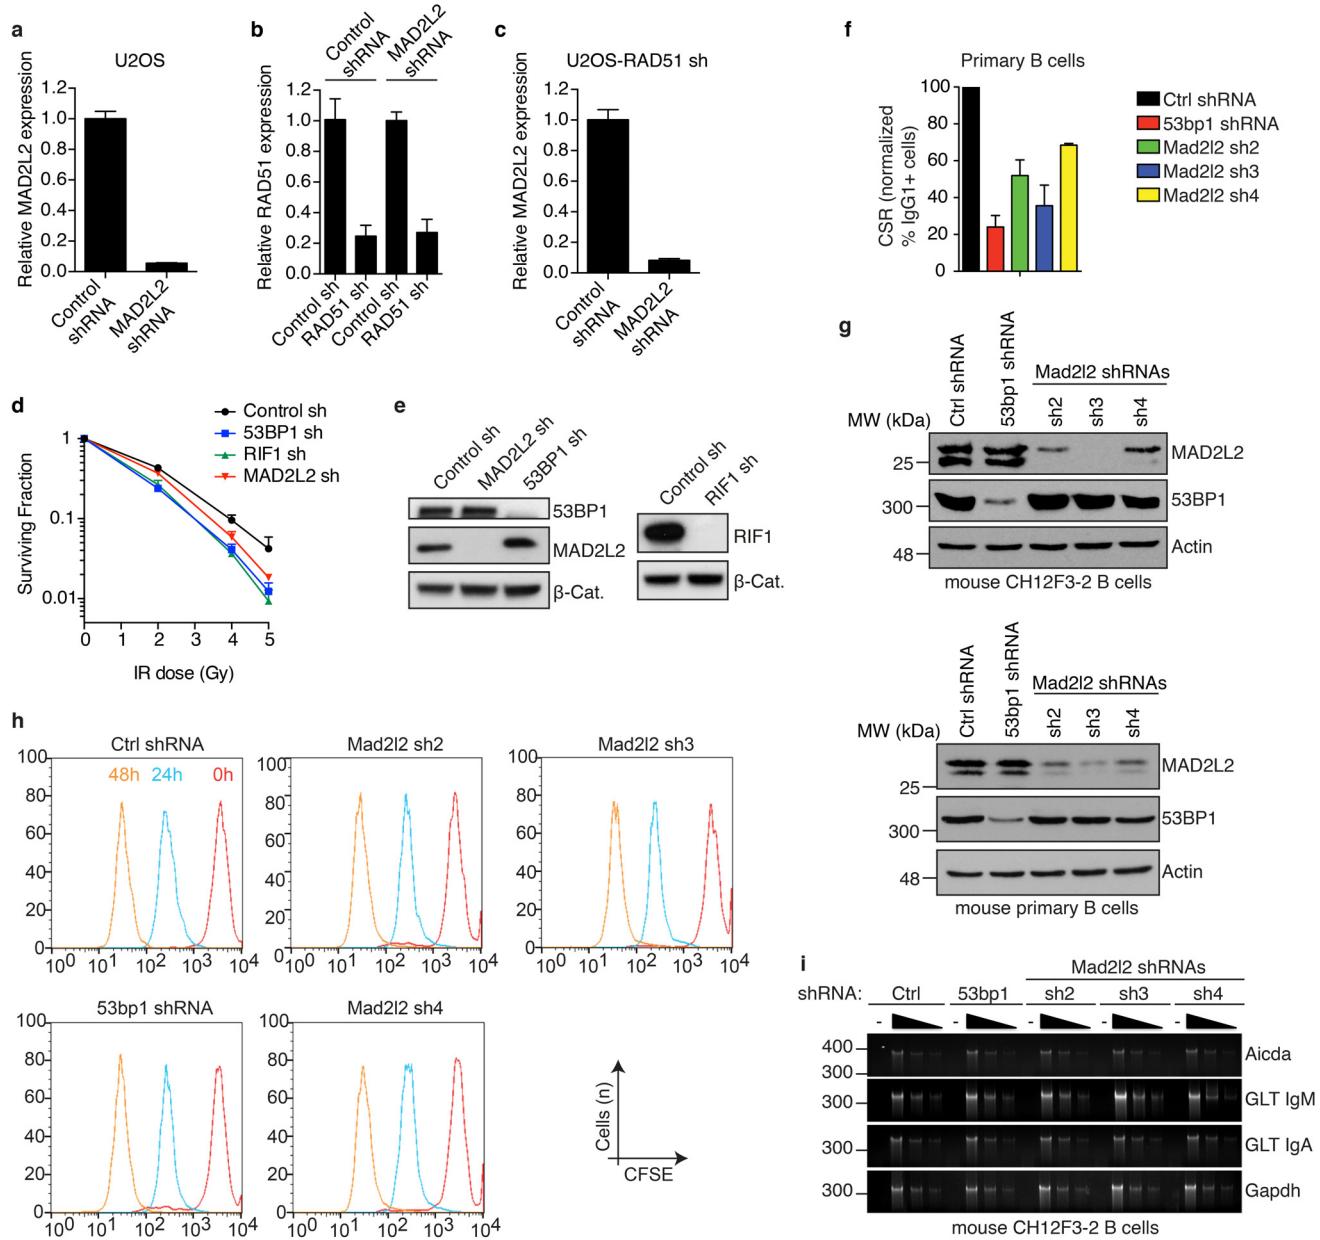
Extended Data Figure 1 | Related to Fig. 1. **a**, A functional genetic screen for TIGIRs identifies independent shRNAs against *Mad2l2* and the previously identified regulators of NHEJ-mediated telomere fusion *Atm*, *Nbs1* (also known as *Nbn*), *Rad50*, *53bp1* and *Rnf8*. Listed are the independent shRNAs enriched >1.5-fold in at least two out of three TIGIR screens, with their average ratio of enrichment over all three screens. Ratios reflect shRNA abundance after 12 days of telomere uncapping at 39 °C followed by 4 days recovery at 32 °C versus shRNA abundance after growth for 4 days at 32 °C. **b**, Additional shRNAs targeting *Mad2l2* or known TIGIRs that were enriched >1.5-fold in one out of three screens, with their ratio. Not shown are additional shRNAs against factors not previously implicated in control of telomere fusion, that were considerably enriched in these TIGIR screens but await validation. **c**, qRT-PCR analysis of *Mad2l2* expression levels in TRF2ts MEFs transduced

with four independent shRNAs targeting *Mad2l2* and used in Fig. 1b (error bars denote s.d.). **d**, Survival assay of TRF2ts cells infected with control or *Mad2l2* shRNAs, complemented with empty control or RNAi-resistant Flag-MAD2L2^{RR} and grown as indicated. **e**, Western blot showing expression of endogenous MAD2L2 and exogenous Flag-MAD2L2 in TRF2ts cells used in **d** and in Fig. 1c. **f**, Photograph of shRNA-transduced TRF2ts cells grown for 12 days at 39 °C. Scale bar, 100 μm. **g**, None out of ten shRNAs targeting *Rev1* or *Rev3* were substantially enriched in any of three independent DDR TIGIR screens. **h**, shRNA-mediated *Rev1* knockdown does not increase survival after prolonged telomere uncapping in survival assays of TRF2ts MEFs. Of note, *Rev3* knockdown compromised viability and was therefore not informative in these assays. **i**, qRT-PCR analysis of mouse *Rev1* expression levels of cells shown in **h** (error bars denote s.d.).



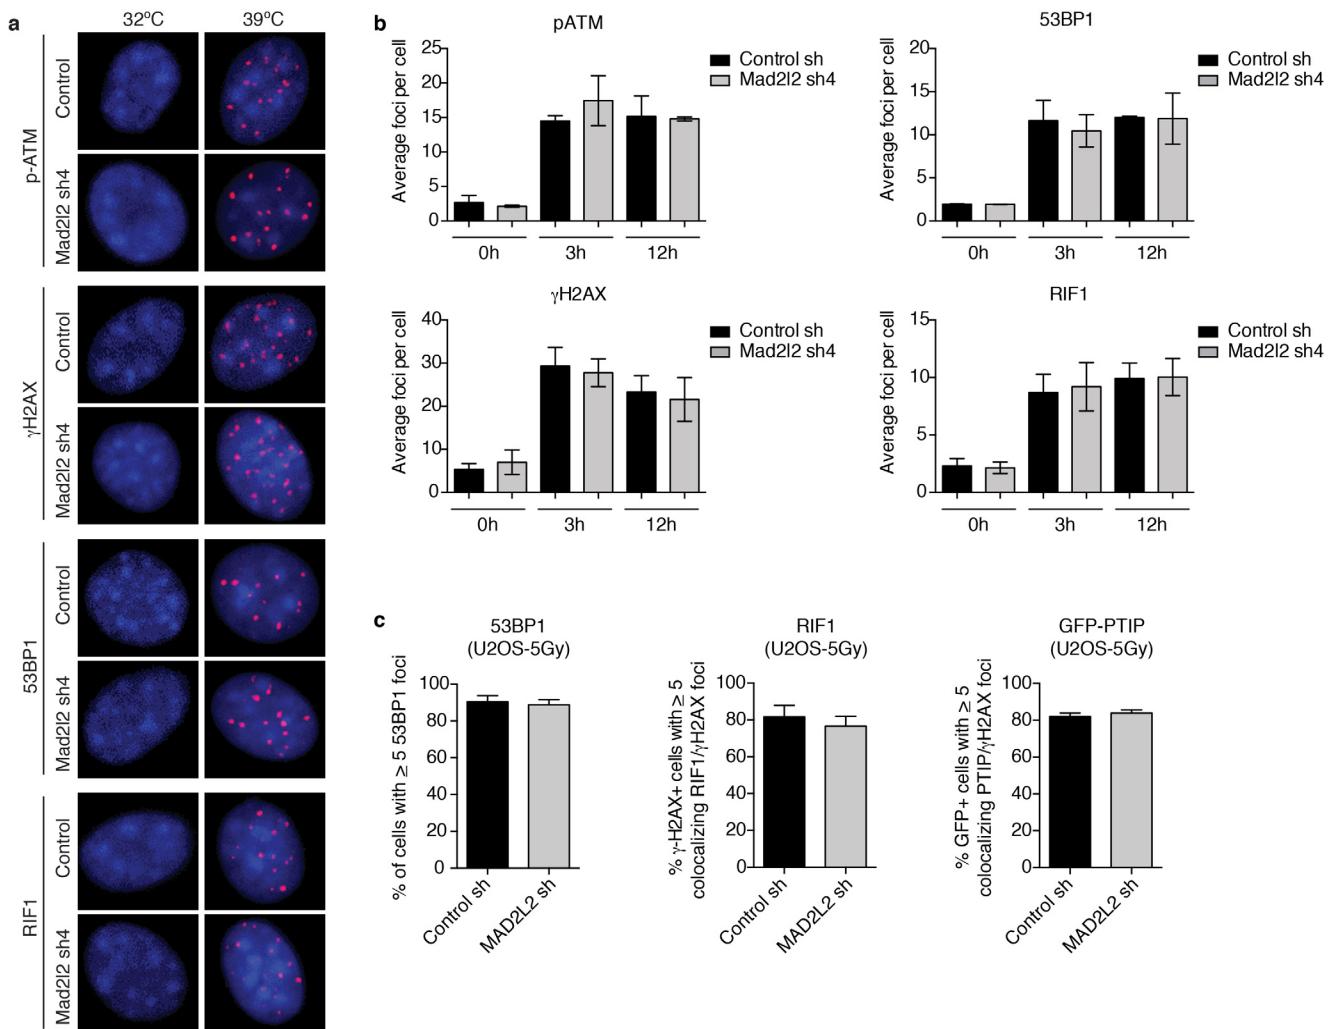
Extended Data Figure 2 | Related to Fig. 2. **a**, Representative metaphase spreads of TRF2ts MEFs transduced as indicated, collected after 24 h at 39 °C for telomere FISH. Original magnifications, $\times 63$. **b**, Representative telomeric single-stranded G-overhang (ss TTAGGG) and total telomere (total TTAGGG) analysis in TRF2ts MEFs at 32 °C and after 48 h at 37 °C or 39 °C. **c, d**, *Mad2l2* knockdown does not affect cell cycle parameters in TRF2ts MEFs. Cell cycle phase analysis was based on propidium iodide staining of asynchronously growing cells (**c**), as well as on 1 h incubation with BrdU, followed by detection of BrdU incorporation and propidium iodide staining for DNA content (**d**) ($n = 3$, mean \pm s.d.).

sister-telomere fusion upon activation of DNA repair in mitosis. Sister-telomere fusions were quantified in IMR90 cells expressing exogenous wild-type 53BP1 and RNF8, or 53BP1 (Thr1608Ala/Ser1618Ala) (TASA) and RNF8(Thr198Ala) (TA) mutant alleles, and depleted for endogenous RNF8 and 53BP1, as well as depleted for MAD2L2, RIF1 or PTIP ($n = 4$). **f**, Examples of DNA content profiles of TRF2ts cells transduced with the indicated shRNAs, grown at 32 °C, or for 48 h at 39 °C and stained with propidium iodide. Analysis of the fraction of cells with 8N (tetraploid) or $>4N$ (aneuploid) DNA content was done on corresponding dotplots of which the results are shown in Fig. 2c.



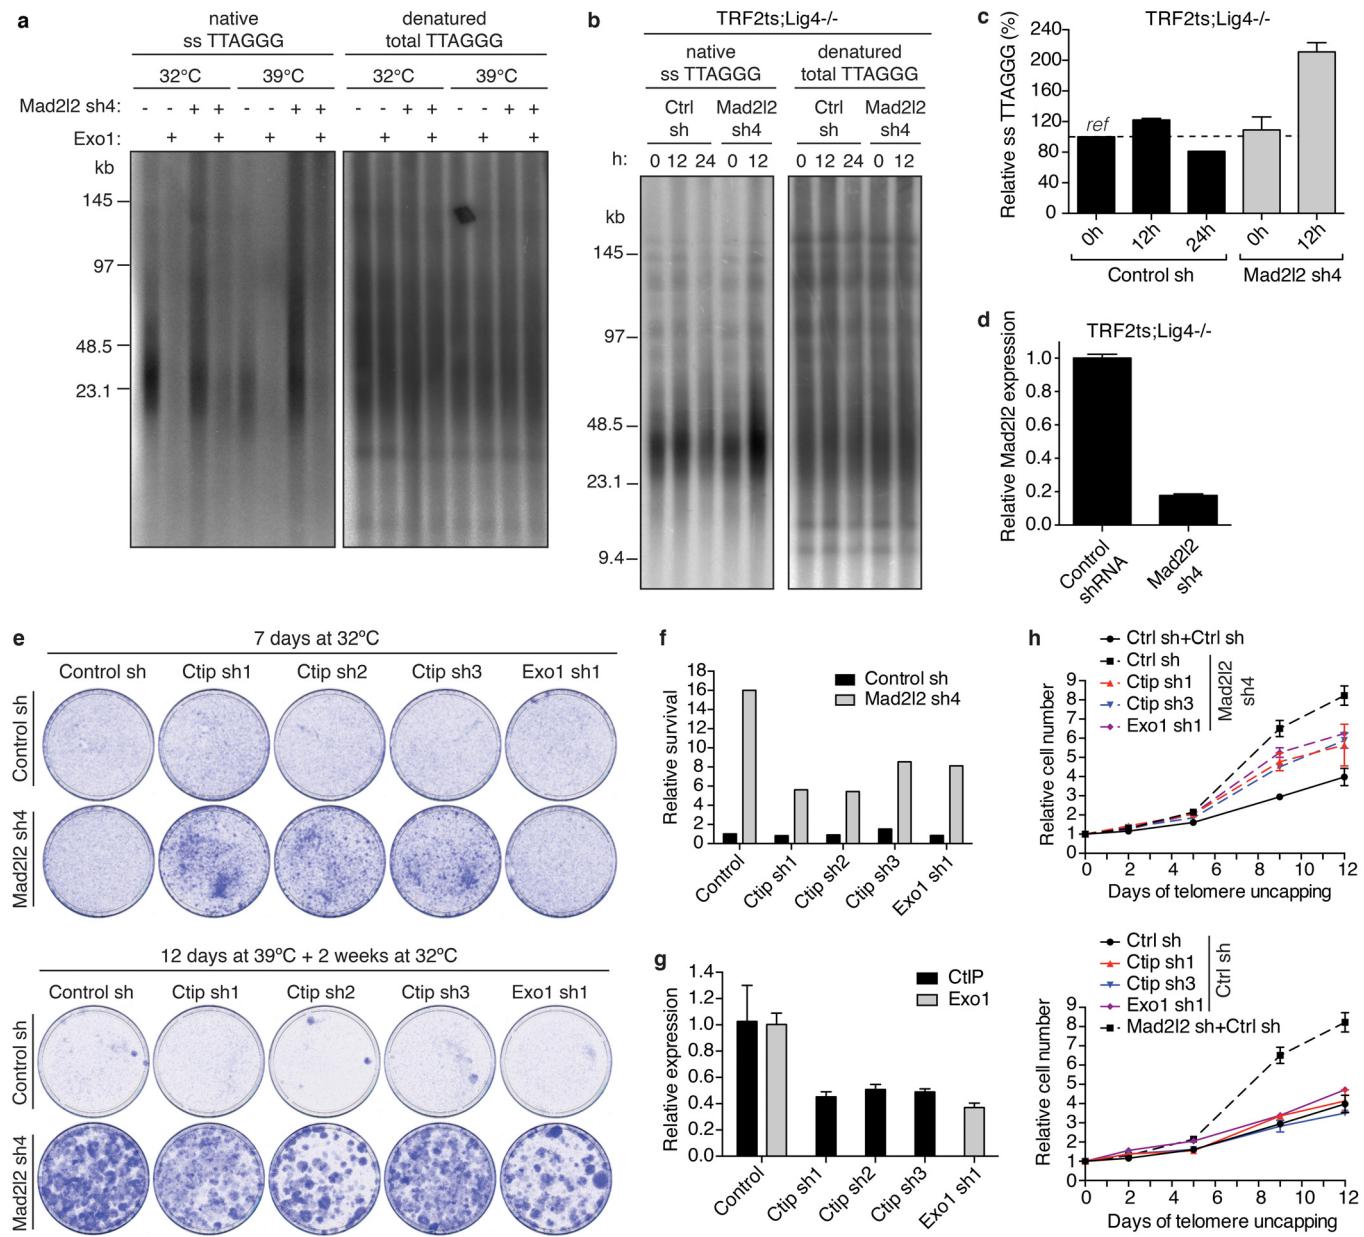
Extended Data Figure 3 | Related to Fig. 2. **a**, qRT-PCR analysis of *MAD2L2* expression levels in U2OS cells infected with control or *MAD2L2* shRNAs, and used in the repair assays shown in Fig. 2d (error bars denote s.d.). **b, c**, qRT-PCR analysis of *RAD51* (**b**) and *MAD2L2* (**c**) expression levels in RAD51-depleted, E6E7-expressing U2OS cells used in the assays shown in Fig. 2e (error bars denote s.d.). **d**, Clonogenic survival assays of U2OS cells transduced with non-targeting control or *53BP1*, *RIF1* or *MAD2L2* shRNAs and treated with the indicated doses of ionizing radiation ($n = 3\text{--}4$, mean \pm s.e.m.). **e**, Western blot analysis of *53BP1*, *MAD2L2* and *RIF1* in U2OS cells transduced with the indicated shRNAs. **f**, CSR in shRNA-transduced primary B cells ($n = 2$, mean \pm s.d.). **g**, Western blot analysis of *MAD2L2* and *53BP1* in

CH12F3-2 B cells and mouse primary B cells transduced with the indicated shRNAs. **h**, *MAD2L2* depletion does not affect cellular proliferation in murine B cells. CH12F3-2 cells transduced with control, *53bp1* or *Mad2l2* shRNAs were loaded with CFSE and analysed by flow cytometry at 0, 24 and 48 h after stimulation. Profiles from all time points are plotted in the same histogram. **i**, *MAD2L2* depletion does not affect the transcription of critical genes implicated in CSR. RT-PCR analysis of *Aid* (*Aicda*) mRNA, IgM and IgA germline transcript (GLT) levels using twofold serial dilutions of cDNA made from activated CH12F3-2 B cells transduced with the indicated shRNAs. *Gapdh* was used as a control for transcript expression.



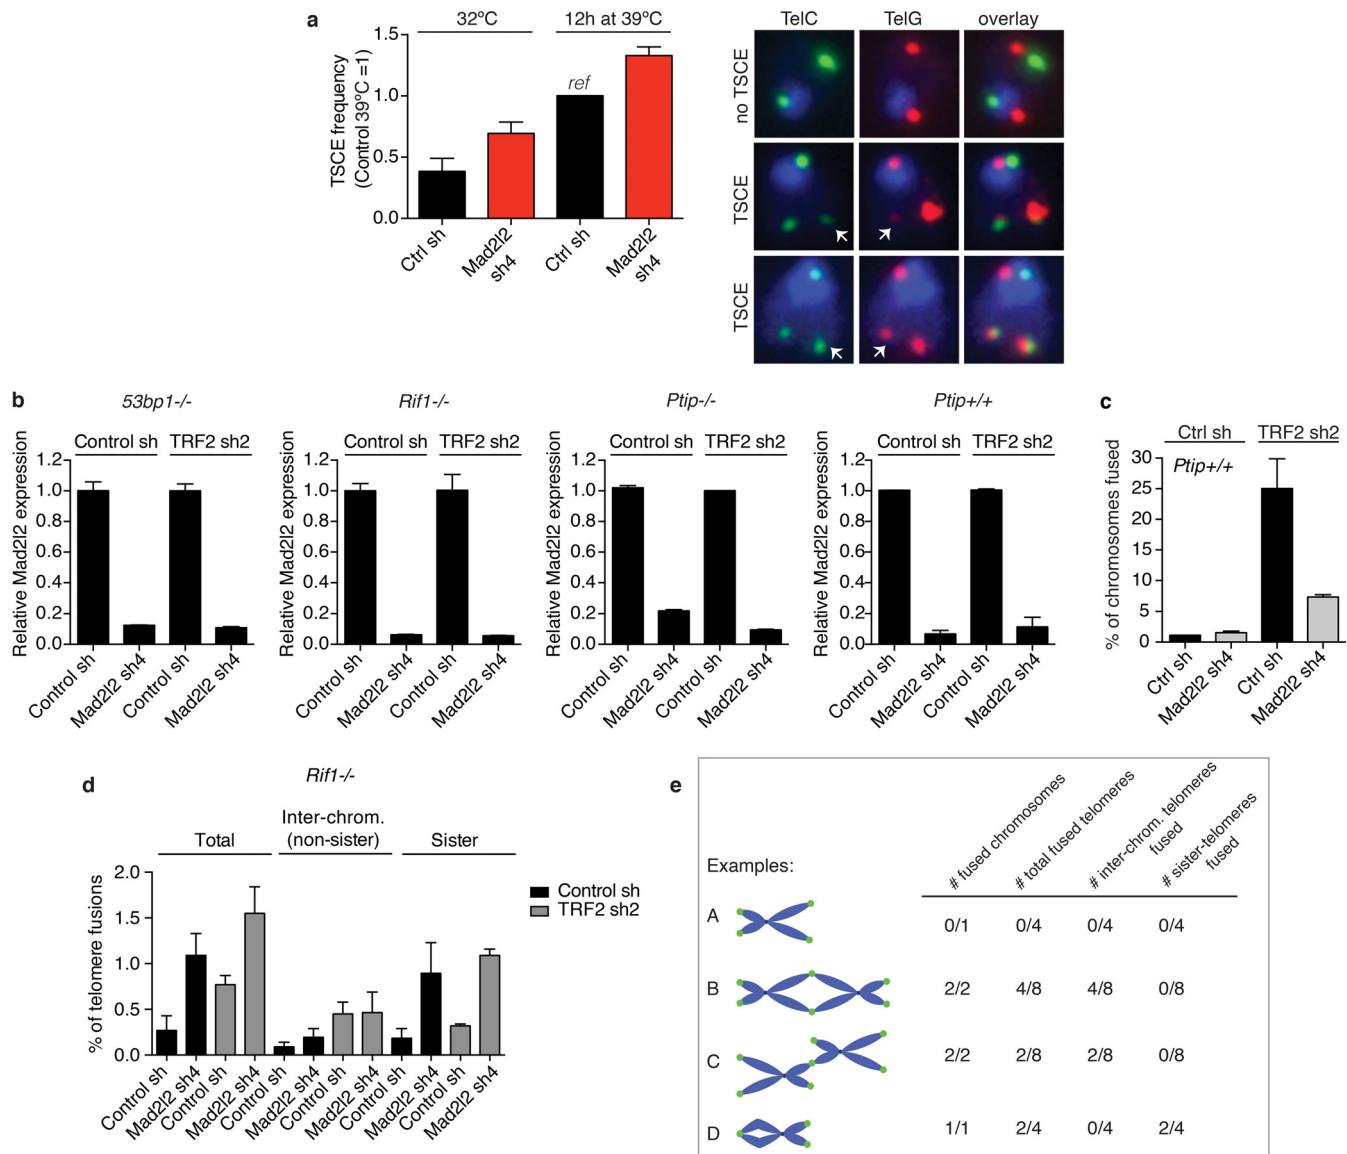
Extended Data Figure 4 | Related to Fig. 3. **a**, Representative images of immunofluorescence detection of p-ATM(S1981), γ H2AX, 53BP1 and RIF1 in TRF2ts MEFs transduced with control or *Mad2l2* shRNAs and grown at 32 °C or for 12 h at 39 °C to induce telomere uncapping. DNA was stained with 4',6-diamidino-2-phenylindole (DAPI). Original magnifications, $\times 40$.

b, Quantification of the number of p-ATM, γ H2AX, 53BP1 and RIF1 foci per cell in TRF2ts MEFs transduced with control or *Mad2l2* shRNAs and grown at 32 °C or for 3 and 12 h at 39 °C ($n = 2$, mean \pm s.d.). **c**, Quantifications of 53BP1, RIF1 and PTIP foci in U2OS cells transduced with control or *MAD2L2* shRNAs, 3 h after 5 Gy ($n = 2$, mean \pm s.d.).



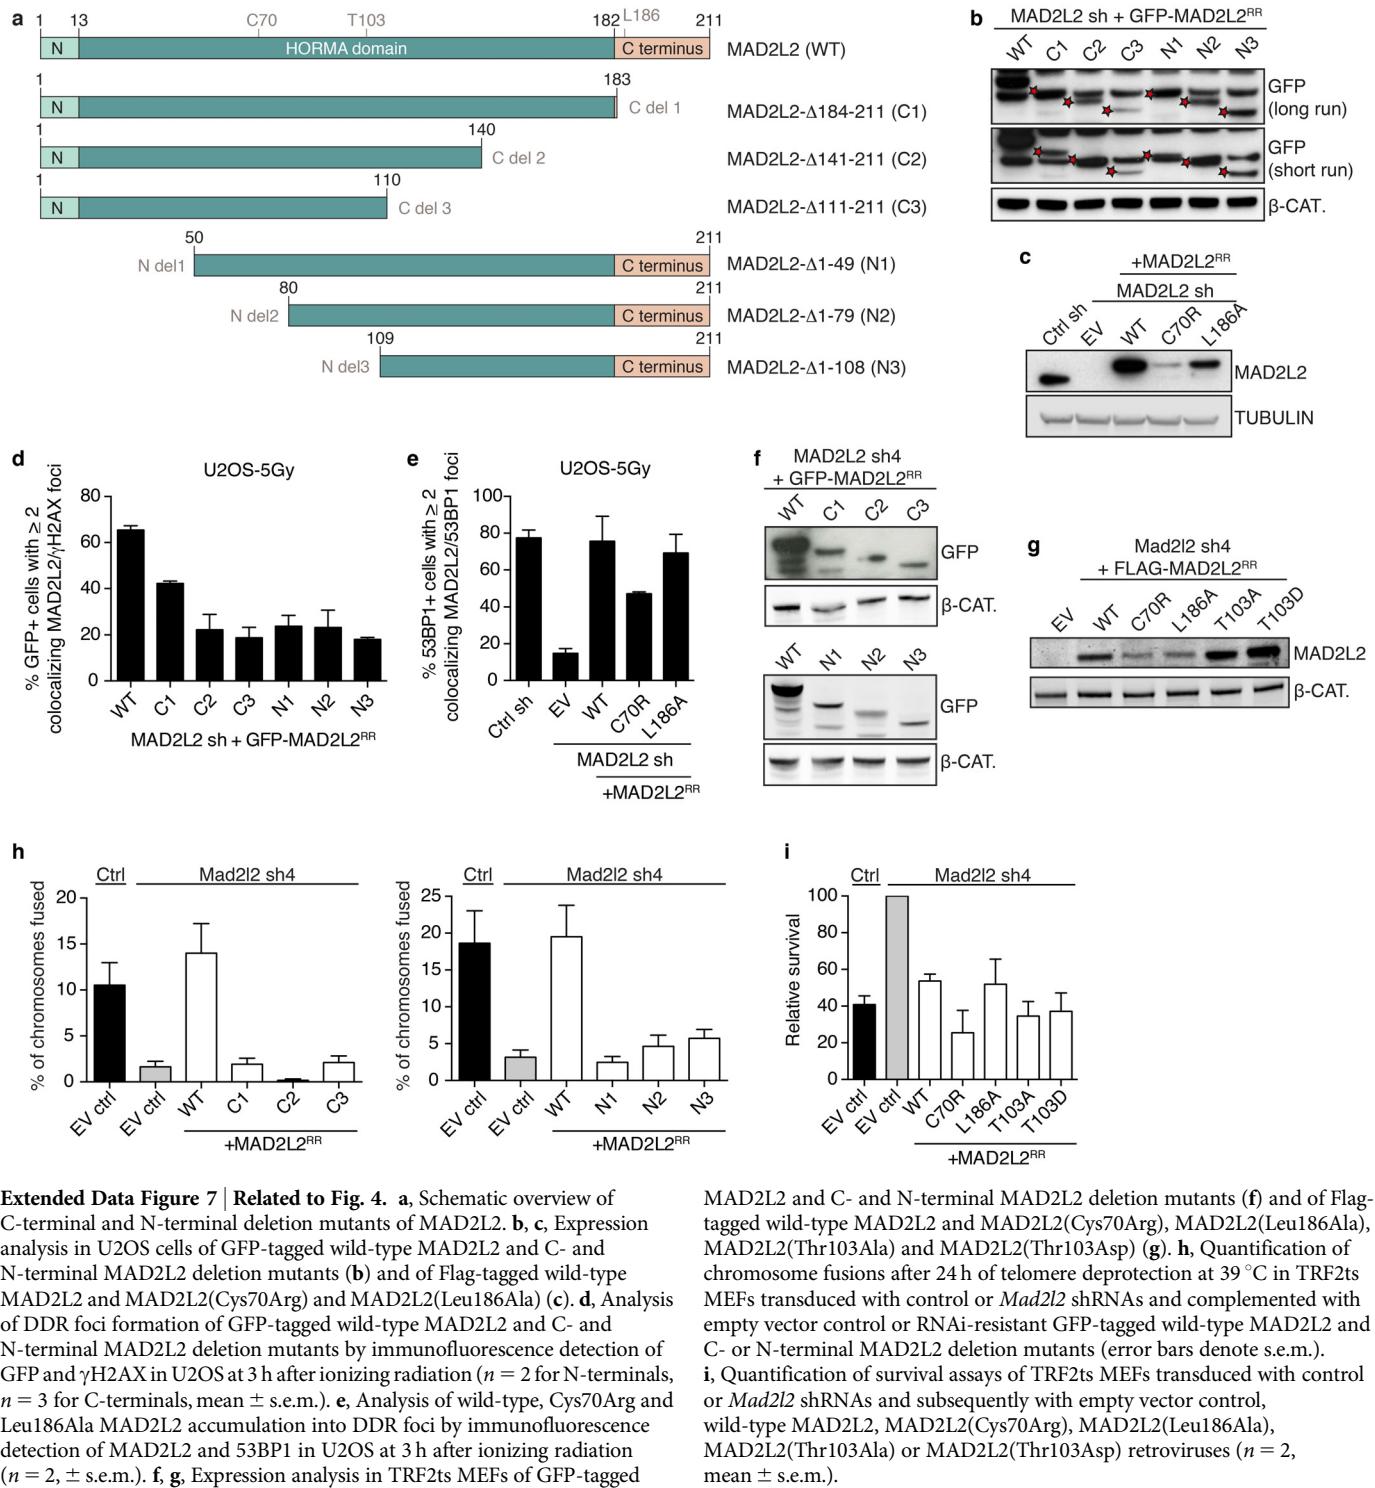
Extended Data Figure 5 | Related to Fig. 3. **a**, Telomeric single-stranded G-overhang assay of TRF2ts MEFs transduced with control or *Mad2l2* shRNA, showing that the increase in overhang signal upon *Mad2l2* knockdown is due to 3' terminal sequences because the signal is removed by treatment with *Escherichia coli* 3' exonuclease Exo1. **b**, *Mad2l2* knockdown causes increased single-stranded telomeric G-overhang signals in TRF2ts *Lig4*^{-/-} MEFs. **c**, Quantification of relative telomeric G-overhang signals in TRF2ts *Lig4*^{-/-} MEFs transduced with control or *Mad2l2* shRNAs and grown at 32 °C or for 12 or 24 h at 39 °C ($n = 2$, mean ± s.e.m.). **d**, qRT-PCR analysis

of *Mad2l2* expression levels in TRF2ts *Lig4*^{-/-} MEFs infected with control or *Mad2l2* shRNA lentivirus (error bars denote s.d.). **e**, Survival assays of TRF2ts MEFs transduced with control or *Mad2l2* shRNAs and subsequently with control, *Ctip* or *Exo1* shRNAs. **f**, Quantification of the survival assays shown in **e**. **g**, qRT-PCR analysis of *Ctip* and *Exo1* expression levels of cells shown in **e**, **f** and **h** and in Fig. 3f (error bars denote s.d.). **h**, Growth curves at 39 °C of TRF2ts MEFs transduced with non-targeting control or *Mad2l2* shRNAs and subsequently with control, *Ctip* or *Exo1* shRNAs (error bars denote s.e.m.).



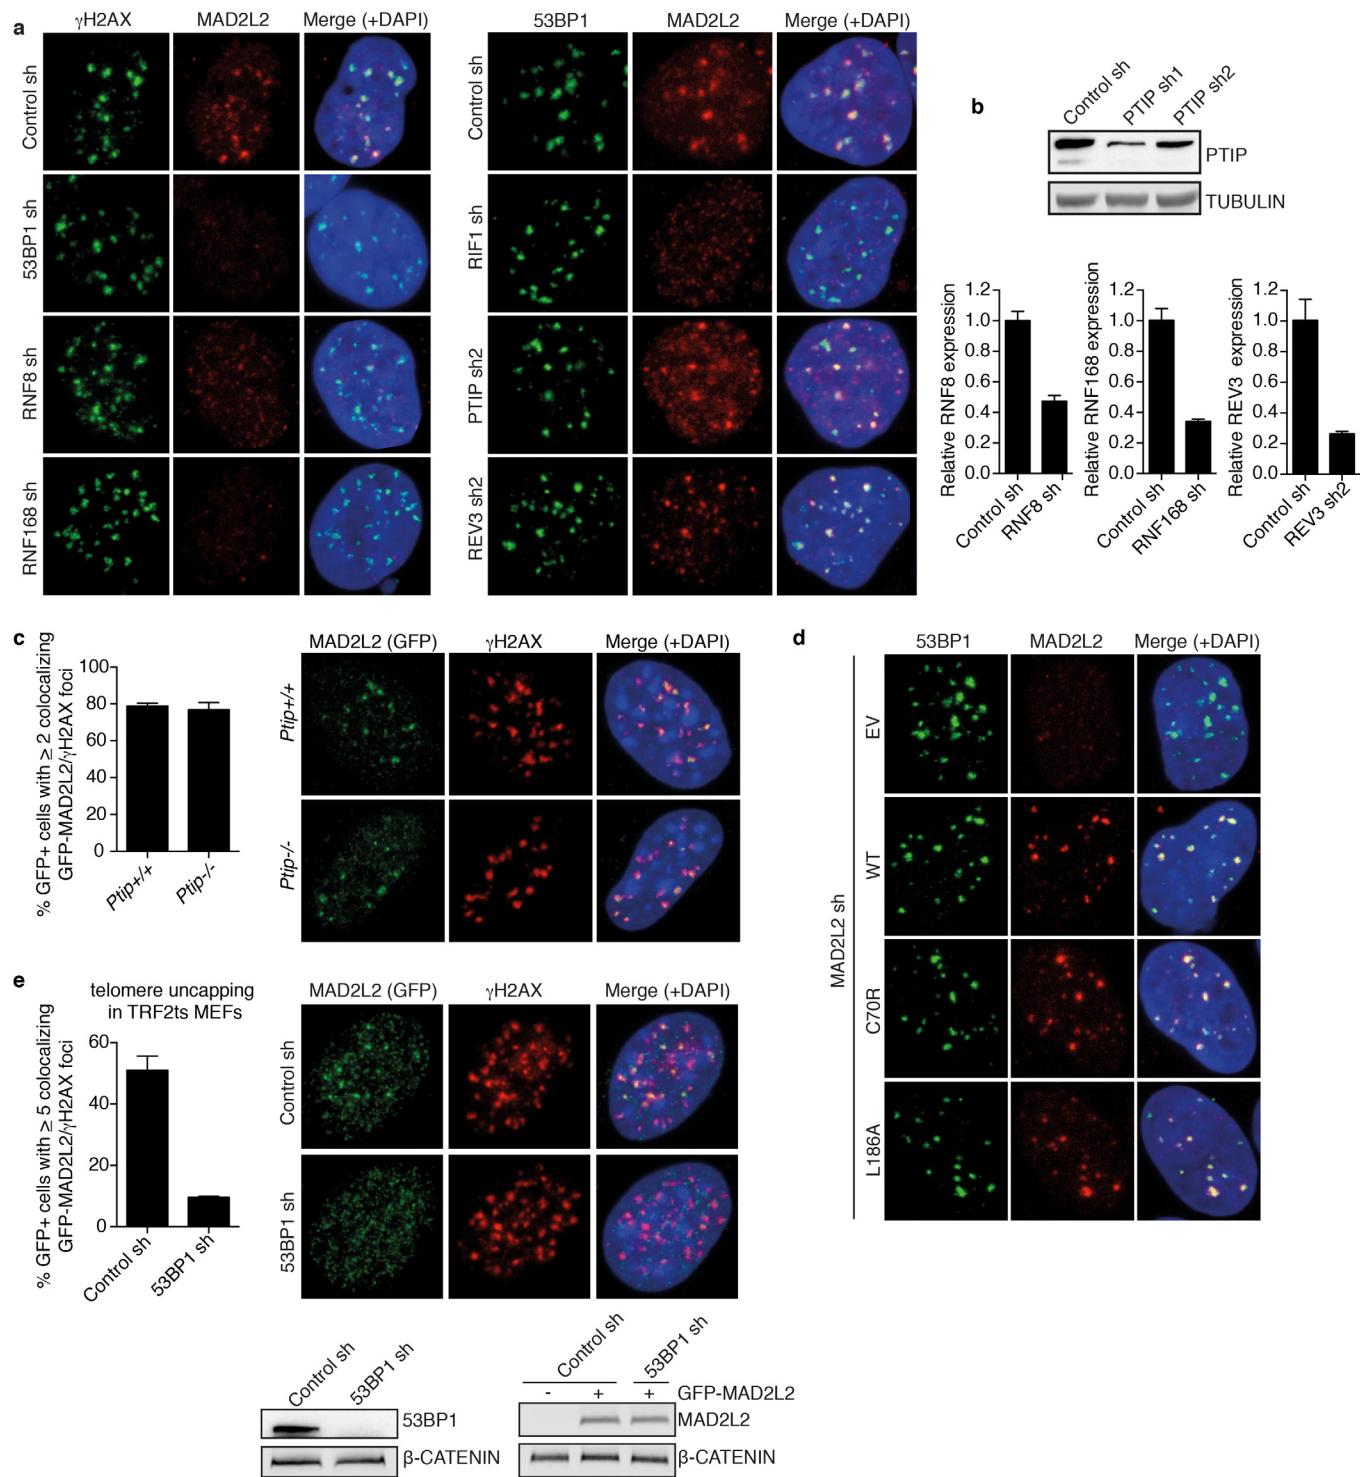
Extended Data Figure 6 | Related to Figs 3 and 4. **a**, TSCE analysis in shRNA-transduced TRF2 MEFs grown at 32 °C or for 12 h at 39 °C to uncap telomeres ($n = 2$, mean \pm s.e.m., counting $>1,000$ chromosomes per condition, per experiment). TSCE frequency in control cells is set at 1 (corresponding to an average of 6.9% of chromosomes with a TSCE event). Shown on the right are examples of chromosomes without and with TSCE in cells quantified on the left. Original magnifications, $\times 63$. **b**, qRT-PCR analysis of *Mad2l2* expression levels in *53bp1*^{-/-}, *Rif1*^{-/-}, *Ptip*^{-/-} and *Ptip*^{+/+} MEFs used in the chromosome fusion analysis shown in Fig. 4a and in c (error bars denote s.d.). **c**, Percentage of chromosomes fused upon TRF2

inhibition in the *Ptip*^{+/+} MEFs matching with the *Ptip*^{-/-} MEFs shown in Fig. 4a ($n = 2$, mean \pm s.e.m.). **d**, Analysis of different types of telomere fusions in *Rif1*^{-/-} MEFs. Depletion of MAD2L2 in *Rif1*^{-/-} MEFs does not reduce inter-chromosomal telomere fusions induced by TRF2 inhibition, indicating epistasis. However, irrespective of TRF2 inhibition, MAD2L2 depletion in *Rif1*^{-/-} MEFs induces association between sister telomeres, causing an increase in total fusions scored for MAD2L2-depleted *Rif1*^{-/-} MEFs, as also visible in Fig. 4a ($n = 2$, mean \pm s.e.m., $>1,300$ – $2,000$ chromosomes were analysed per condition, per independent experiment). **e**, Explanation of scoring different types of telomere fusions shown in d.



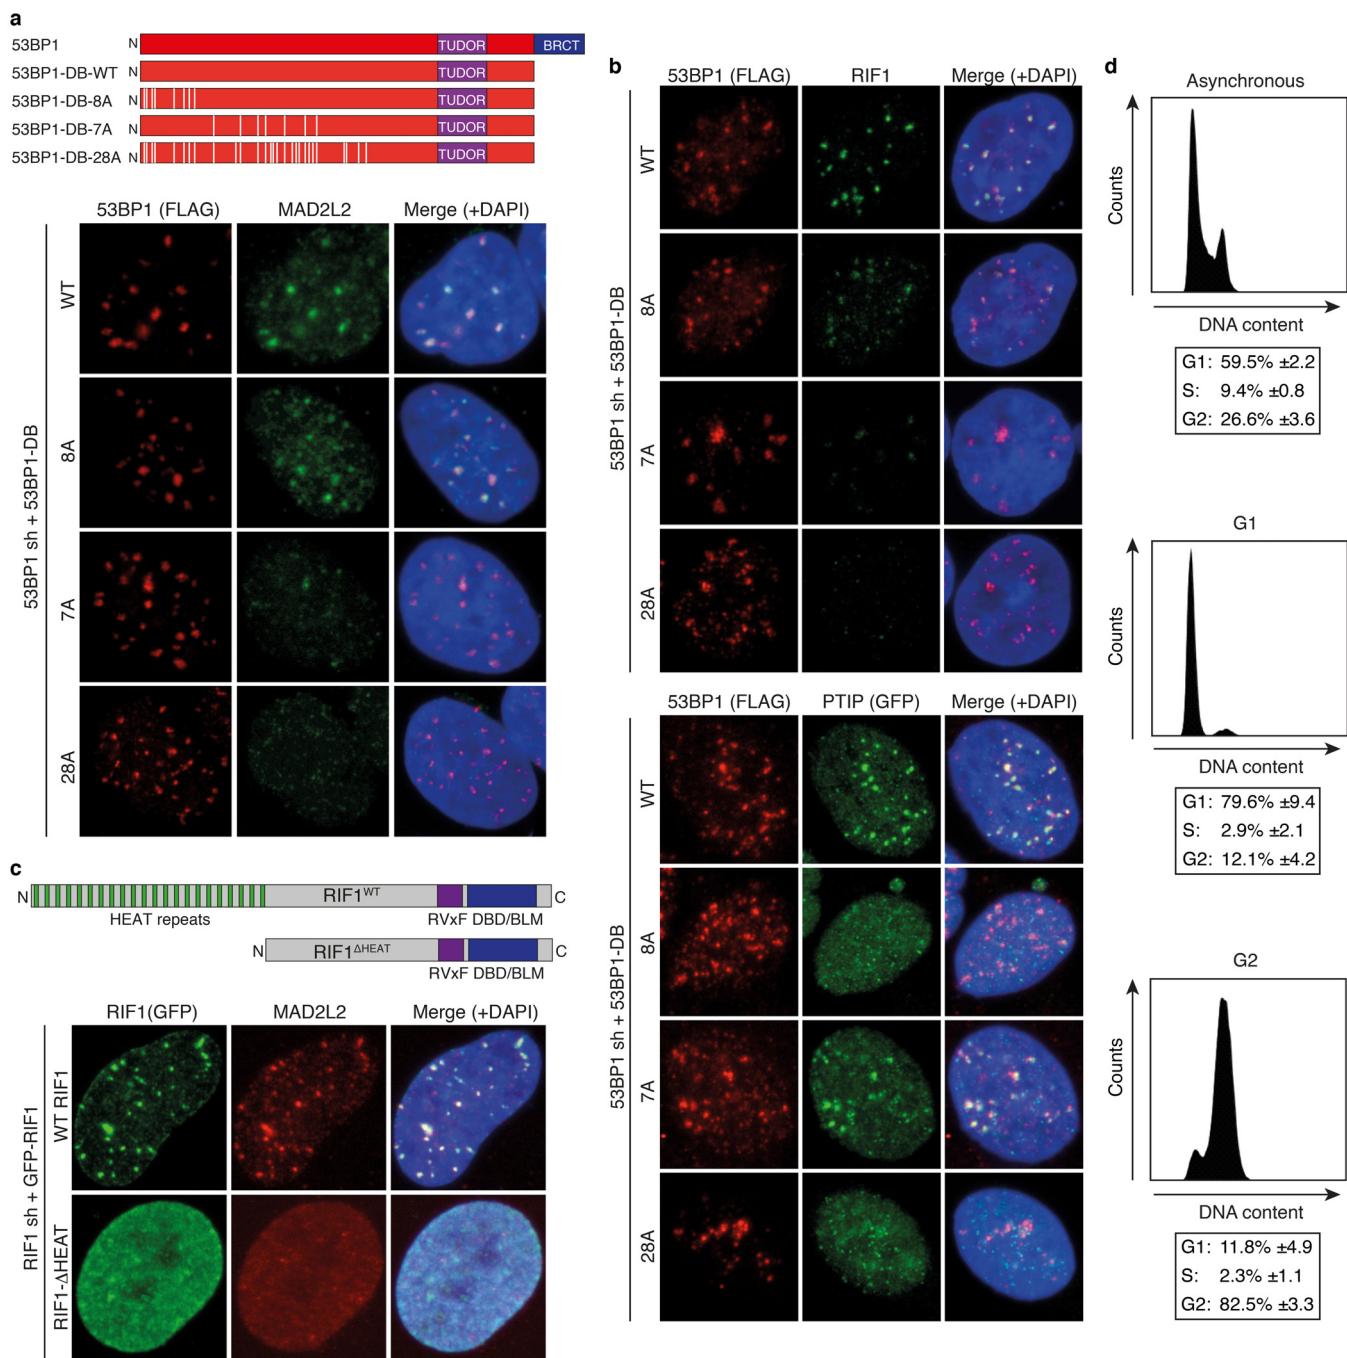
Extended Data Figure 7 | Related to Fig. 4. **a**, Schematic overview of C-terminal and N-terminal deletion mutants of MAD2L2. **b, c**, Expression analysis in U2OS cells of GFP-tagged wild-type MAD2L2 and C- and N-terminal MAD2L2 deletion mutants (**b**) and of Flag-tagged wild-type MAD2L2 and MAD2L2(Cys70Arg) and MAD2L2(Leu186Ala) (**c**). **d**, Analysis of DDR foci formation of GFP-tagged wild-type MAD2L2 and C- and N-terminal MAD2L2 deletion mutants by immunofluorescence detection of GFP and γ H2AX in U2OS at 3 h after ionizing radiation ($n = 2$ for N-terminals, $n = 3$ for C-terminals, mean \pm s.e.m.). **e**, Analysis of wild-type, Cys70Arg and Leu186Ala MAD2L2 accumulation into DDR foci by immunofluorescence detection of MAD2L2 and 53BP1 in U2OS at 3 h after ionizing radiation ($n = 2$, \pm s.e.m.). **f, g**, Expression analysis in TRF2ts MEFs of GFP-tagged

MAD2L2 and C- and N-terminal MAD2L2 deletion mutants (**f**) and of Flag-tagged wild-type MAD2L2 and MAD2L2(Cys70Arg), MAD2L2(Leu186Ala), MAD2L2(Thr103Ala) and MAD2L2(Thr103Asp) (**g**). **h**, Quantification of chromosome fusions after 24 h of telomere deprotection at 39 °C in TRF2ts MEFs transduced with control or *Mad2l2* shRNAs and complemented with empty vector control or RNAi-resistant GFP-tagged wild-type MAD2L2 and C- or N-terminal MAD2L2 deletion mutants (error bars denote s.e.m.). **i**, Quantification of survival assays of TRF2ts MEFs transduced with control or *Mad2l2* shRNAs and subsequently with empty vector control, wild-type MAD2L2, MAD2L2(Cys70Arg), MAD2L2(Leu186Ala), MAD2L2(Thr103Ala) or MAD2L2(Thr103Asp) retroviruses ($n = 2$, mean \pm s.e.m.).



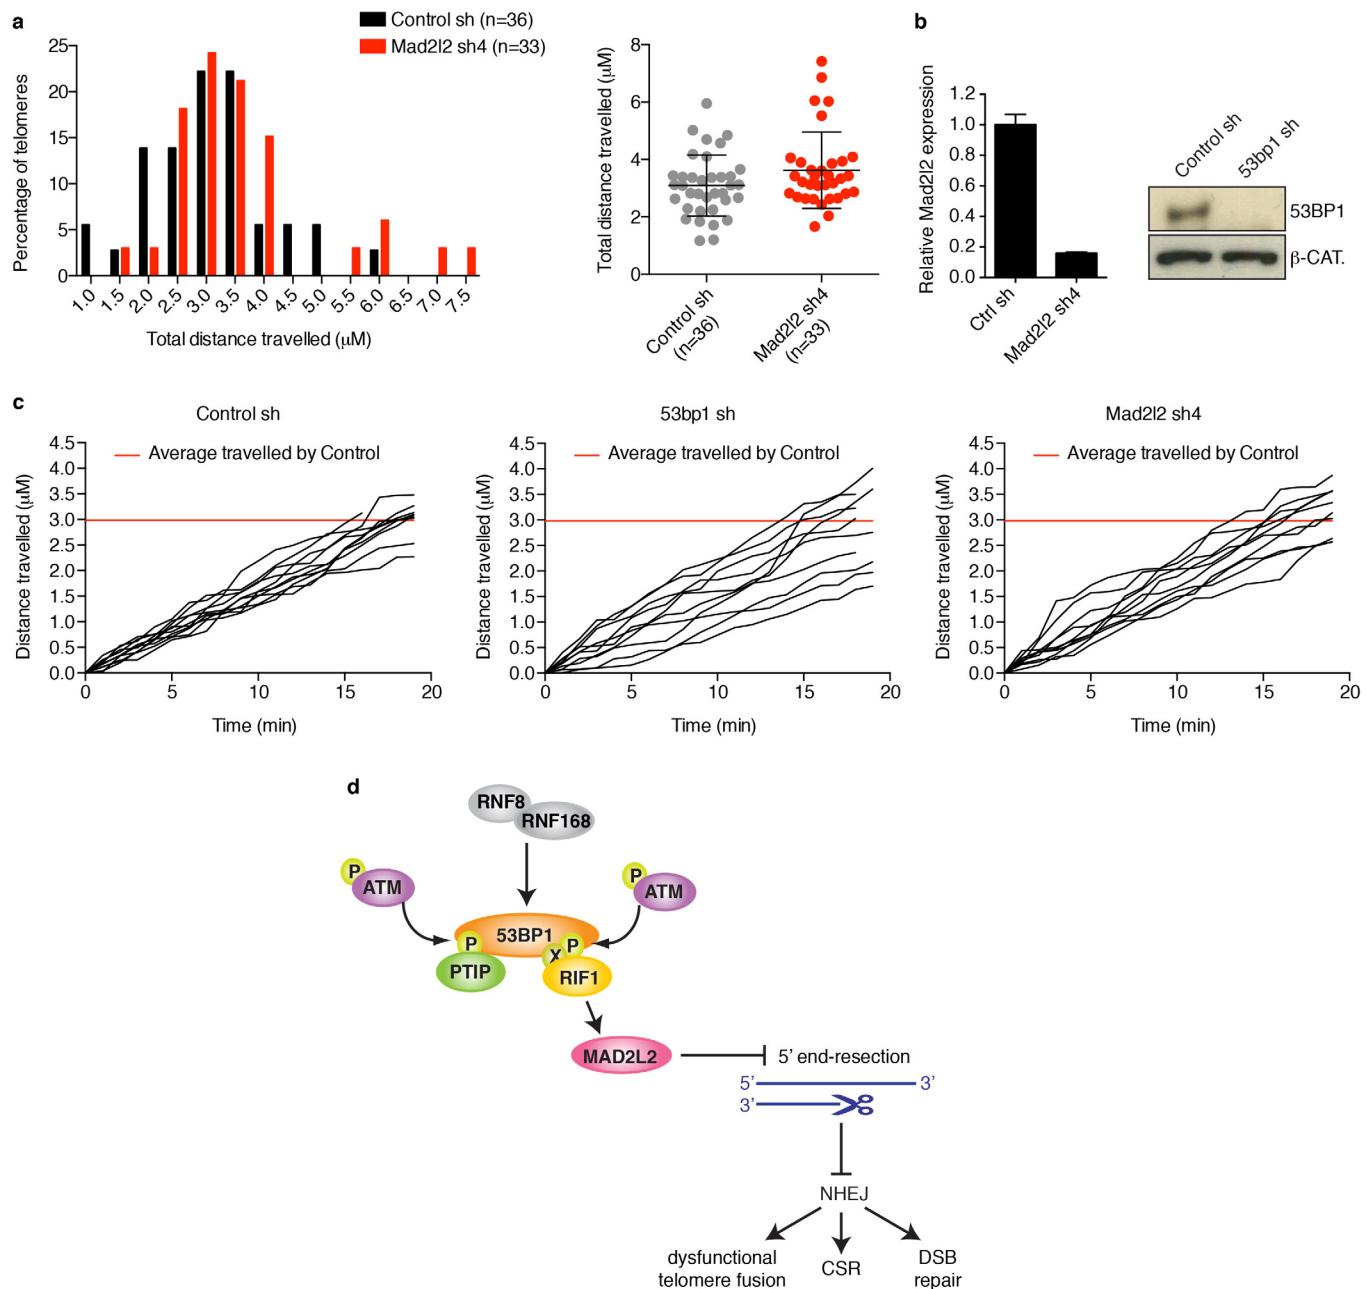
Extended Data Figure 8 | Related to Fig. 4. **a**, Representative images of immunofluorescence detection of endogenous MAD2L2 and γ H2AX or 53BP1 in U2OS cells transduced with control, 53BP1, RNF8, RNF168, RIF1, PTIP or REV3 shRNA lentiviruses, irradiated with 5 Gy and processed for immunofluorescence after 3 h (quantifications are shown in Fig. 4e). **b**, Western blot or qRT-PCR analysis of PTIP, RNF8, RNF168 and REV3 levels in shRNA-transduced U2OS cells (error bars denote s.d.). **c**, Quantification and representative images of immunofluorescence detection of γ H2AX and GFP-MAD2L2 in *Ptip*^{+/+} or *Ptip*^{-/-} MEFs, 3 h after 5 Gy ($n = 2$, mean \pm s.d.). Original magnifications, $\times 63$ with zoom factor 5.

mean \pm s.d.). **d**, Representative images of immunofluorescence for 53BP1 and exogenous Flag-MAD2L2 wild-type, Cys70Arg or Leu186Ala mutants in U2OS cells depleted for endogenous MAD2L2 with lentiviral shRNA, processed for immunofluorescence 3 h after 5 Gy (quantifications are shown in Extended Data Fig. 7e). **e**, Quantification and representative immunofluorescence images of GFP-MAD2L2 localization to uncapped telomeres in TRF2ts MEFs transduced with control or 53bp1 shRNAs ($n = 2$, mean \pm s.d.). Original magnifications, $\times 63$ with zoom factor 5.



Extended Data Figure 9 | Related to Fig. 4. **a**, Schematic representation of the 53BP1 alleles with wild-type and substituted S/TQ sites used to address MAD2L2 IRIF localization dependence on 53BP1. Representative immunofluorescence images are displayed that show colocalization of endogenous MAD2L2 with 53BP1-DB-WT and -8A, but not with 53BP1-DB-7A or -28A. Quantifications are provided in Fig. 4f. **b**, Top, representative immunofluorescence images showing colocalization of endogenous RIF1 with 53BP1-DB-WT and -8A, but not 53BP1-DB-7A or -28A. Quantifications are presented in Fig. 4f. Bottom, representative immunofluorescence images showing colocalization of PTIP-GFP with 53BP1-DB-WT and -7A, but

impaired colocalization of PTIP-GFP with 53BP1-DB-8A or -28A. **c**, Schematic representation of GFP-tagged wild-type RIF1 and GFP-tagged RIF1 lacking the N-terminal HEAT repeats (Δ HEAT), used to address MAD2L2 IRIF localization dependence on the HEAT repeats of RIF1. Representative immunofluorescence images are shown. Quantifications are presented in Fig. 4g. **d**, Cell cycle phase distributions of the RPE cells used in Fig. 4h to address cell cycle dependence of endogenous RIF1 and MAD2L2 localization to IRIFs ($n = 3$, mean \pm s.d.). Original magnifications, $\times 63$ with zoom factor 5.



Extended Data Figure 10 | Related to Fig. 4. **a**, Frequency distribution (left) and scatter plot (right) of total distances travelled by uncapped telomeres in control or *Mad2l2* knockdown cells. **b**, qRT-PCR analysis of *Mad2l2* expression levels and western blot analysis of 53BP1 protein levels in TRF2ts cells used in the experiments shown in **a** and **c** (error bars denote \pm s.d.).

c, Distance travelled by 10 representative uncapped telomeres for each condition. While multiple uncapped telomeres in 53BP1-depleted cells have reduced mobility, this is not seen for uncapped telomeres in MAD2L2-depleted cells. **d**, Model of the role of MAD2L2 in promoting NHEJ.

REV7 counteracts DNA double-strand break resection and affects PARP inhibition

Guotai Xu¹, J. Ross Chapman^{2*}, Inger Brandsma^{3*}, Jingsong Yuan⁴, Martin Mistrik⁵, Peter Bouwman⁶, Jirina Bartkova⁷, Ewa Gogola¹, Daniël Warmerdam⁸, Marco Barazas¹, Janneke E. Jaspers¹, Kenji Watanabe⁷, Mark Pieterse⁶, Ariena Kersbergen¹, Wendy Sol¹, Patrick H. N. Celie⁹, Philip C. Schouten⁶, Bram van den Broek⁸, Ahmed Salman², Marja Nieuwland¹⁰, Iris de Rink¹⁰, Jorma de Ronde¹¹, Kees Jalink⁸, Simon J. Boulton¹², Junjie Chen⁴, Dik C. van Gent³, Jiri Bartek^{5,7}, Jos Jonkers⁶, Piet Borst¹ & Sven Rottenberg^{1,13}

Error-free repair of DNA double-strand breaks (DSBs) is achieved by homologous recombination (HR), and BRCA1 is an important factor for this repair pathway¹. In the absence of BRCA1-mediated HR, the administration of PARP inhibitors induces synthetic lethality of tumour cells of patients with breast or ovarian cancers^{2,3}. Despite the benefit of this tailored therapy, drug resistance can occur by HR restoration⁴. Genetic reversion of BRCA1-inactivating mutations can be the underlying mechanism of drug resistance, but this does not explain resistance in all cases⁵. In particular, little is known about BRCA1-independent restoration of HR. Here we show that loss of REV7 (also known as MAD2L2) in mouse and human cell lines re-establishes CTIP-dependent end resection of DSBs in BRCA1-deficient cells, leading to HR restoration and PARP inhibitor resistance, which is reversed by ATM kinase inhibition. REV7 is recruited to DSBs in a manner dependent on the H2AX-MDC1-RNF8-RNF168-53BP1 chromatin pathway, and seems to block HR and promote end joining in addition to its regulatory role in DNA damage tolerance⁶. Finally, we establish that REV7 blocks DSB resection to promote non-homologous end-joining during immunoglobulin class switch recombination. Our results reveal an unexpected crucial function of REV7 downstream of 53BP1 in coordinating pathological DSB repair pathway choices in BRCA1-deficient cells.

To identify mechanisms of BRCA1-independent restoration of the HR pathway, we carried out a loss-of-function short hairpin RNA (shRNA) screen using the KB1P-B11 and KB1P-G3 cell lines that we previously derived from *Brcal*^{-/-} *p53*^{-/-} (*p53* is also known as *Trp53*) mouse mammary tumours⁷ (Fig. 1a and Supplementary Table 1). Resistant cells were selected with a high concentration of olaparib (500 nM, about 100-fold the half-maximal inhibitory concentration (IC₅₀)), which killed cells of the empty vector control. Sequencing of the olaparib-surviving colonies revealed a reproducible enrichment of various individual hairpins targeting *Rev7* or *53bp1* (also known as *Trp53bp1*). To validate the *Rev7* hit, we introduced two different hairpins into the KB1P-B11 and KB1P-G3 cell lines; these substantially inhibited *Rev7* expression (Fig. 1b, c and Extended Data Fig. 1a). Despite the role of REV7 in metaphase-to-anaphase transition⁸, the level of *Rev7* inhibition in these cells did not affect proliferation (Extended Data Fig. 1b, c), allowing long-term clonogenic survival assays. We confirmed that loss of *Rev7* resulted in increased resistance to the PARP inhibitors (PARPi) olaparib and AZD2461 (ref. 7) in both cell lines (Fig. 1d and Extended Data

Fig. 1d–g). Resistant cells that survived olaparib treatment (*Rev7* sh1/2-ola) yielded even lower REV7 expression levels and increased numbers of colonies after PARPi treatment (Fig. 1b–d and Extended Data Fig. 1h). When we reconstituted the *Rev7*-depleted cells with shRNA-resistant *Rev7* complementary DNA resulting in similar REV7 protein levels (Extended Data Fig. 1i), we successfully re-sensitized the tumour cells to PARPi (Fig. 1e, f).

Tumours derived from the *Brcal*^{-/-} *p53*^{-/-} cells with stable *Rev7* inhibition also showed olaparib resistance *in vivo*, in contrast to the empty vector controls (Fig. 1g and Extended Data Fig. 1j–l). In addition, we found that *Rev7* loss explains some cases of *in vivo* acquired PARPi resistance in BRCA1-deficient mouse mammary tumours (data not shown). *REV7* depletion also resulted in PARPi resistance of the human BRCA1-deficient cell line SUM149PT (Extended Data Fig. 2). Together, these data strongly indicate that inhibition of *Rev7* confers PARPi resistance in BRCA1-deficient tumour cells.

Together with the catalytic subunit REV3, REV7 forms the translation synthesis polymerase ζ (Pol ζ), and it interacts with REV1 (ref. 9). We therefore investigated whether loss of REV1 or REV3 also confers PARPi resistance in *Brcal*^{-/-} *p53*^{-/-} cells. A 60% inhibition of *Rev1* or *Rev3* transcripts did not cause olaparib resistance (Extended Data Fig. 3a–d). Moreover, we studied various shRNA-resistant REV7 mutants that lack REV1 (Leu186Ala/Gln200Ala/Tyr202Ala and a 1–183-amino-acids truncated protein) or REV3 (Cys70Arg) binding sites^{10,11}. In contrast to the truncated 1–140-amino-acid REV7 protein, these mutants are recruited to DNA damage sites (Extended Data Fig. 3e–g), and their expression in *Rev7* shRNA KB1P-B11 and KB1P-G3 cells significantly restored the sensitivity to PARPi to a degree approaching that of wild-type REV7 (Fig. 2a, b; $P < 0.001$, *t*-test). The remaining differences of the Leu186Ala/Gln200Ala/Tyr202Ala or Cys70Arg mutants with wild-type REV7 may be explained by unequal expression levels (Extended Data Fig. 3f). These data suggest that the REV1 or REV3 interaction is not absolutely required for the REV7-mediated function in this context.

We observed that green fluorescent protein (GFP)-tagged REV7 colocalizes with 53BP1 shortly after DNA damage induction, suggesting that REV7 acts directly at the site of DNA damage (Fig. 2c). REV7 recruitment depends on H2AX, MDC1, RNF8, RNF168 and partly ATM, in both mouse and human cells (Fig. 2d, e and Extended Data Fig. 4a–d). To examine whether PARPi resistance in *Rev7*-depleted *Brcal*^{-/-} *p53*^{-/-} tumour cells is due to HR restoration, we investigated RAD51 focus

¹Division of Molecular Oncology, The Netherlands Cancer Institute, Plesmanlaan 121, 1066CX Amsterdam, The Netherlands. ²The Wellcome Trust Centre for Human Genetics, Roosevelt Drive, Oxford OX3 7BN, UK. ³Department of Genetics, Erasmus University Medical Center, 3000 CA Rotterdam, The Netherlands. ⁴Department of Experimental Radiation Oncology, University of Texas MD Anderson Cancer Center, Houston, Texas 77030, USA. ⁵Institute of Molecular and Translational Medicine, Faculty of Medicine and Dentistry, Palacky University, 779 00 Olomouc, Czech Republic. ⁶Division of Molecular Pathology, The Netherlands Cancer Institute, Plesmanlaan 121, 1066CX Amsterdam, The Netherlands. ⁷Danish Cancer Society Research Center, 2100 Copenhagen, Denmark. ⁸Division of Cell Biology, The Netherlands Cancer Institute, Plesmanlaan 121, 1066CX Amsterdam, The Netherlands. ⁹Protein Facility, The Netherlands Cancer Institute, Plesmanlaan 121, 1066CX Amsterdam, The Netherlands. ¹⁰Deep Sequencing Core Facility, The Netherlands Cancer Institute, Plesmanlaan 121, 1066CX Amsterdam, The Netherlands. ¹¹Division of Molecular Carcinogenesis, The Netherlands Cancer Institute, Plesmanlaan 121, 1066CX Amsterdam, The Netherlands. ¹²DNA Damage Response Laboratory, London Research Institute, Cancer Research UK, Clare Hall, South Mimms, Hertfordshire EN6 3LD, UK.

¹³Institute of Animal Pathology, Vetsuisse Faculty, University of Bern, Laenggassestrasse 122, 3012 Bern, Switzerland.
*These authors contributed equally to this work.

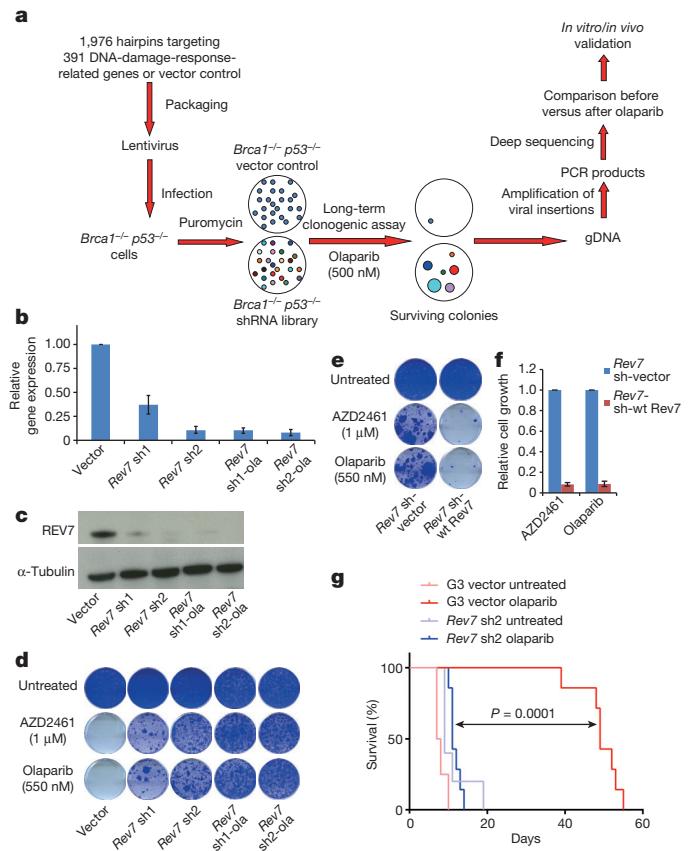


Figure 1 | Identification of loss of Rev7 in PARPi-resistant *Brca1*^{-/-} *p53*^{-/-} mammary tumour cells. **a**, Design of the functional shRNA screen. gDNA, genomic DNA. **b, c**, Quantification of *Rev7* transcript (**b**) or protein (**c**) levels in KB1P-G3 cells transduced with *Rev7*-targeting shRNAs or the vector control. *Hprt* was used as a control for transcript expression, and β-tubulin was used as a control for protein expression. The data represent the mean ± s.d. **d, e**, Long-term clonogenic assay using KB1P-G3 cells transduced with the indicated constructs (wt *Rev7* stands for pLenti6-wt *Rev7*) and treatments. **f**, Quantification of the clonogenic assay in **e** by determining the absorbance of crystal violet at 590 nm. All the groups were normalized to the absorbance of the vector control. The data represent the mean ± s.d. **g**, Overall survival of mice with KB1P-G3-derived *Rev7*-depleted or control tumours treated with one regimen of 50 mg olaparib per kilogram daily for 28 days or left untreated. The *P* value was calculated using the log-rank test.

formation after 10 Gy ionizing radiation. As shown in Fig. 3a, b and Extended Data Fig. 4e, f, *Rev7* loss resulted in the restoration of RAD51 foci formed after DNA damage. To exclude potential off-target effects of the hairpins, we reconstituted *Rev7* shRNA1 and shRNA2 cells with shRNA-resistant mouse or human REV7-GFP fusion proteins (Extended Data Fig. 4g). REV7 re-expression abolished RAD51 focus formation after DNA damage in GFP-positive cells (Fig. 3b). We confirmed the re-appearance of RAD51 foci after tumour irradiation *in vivo* using computed tomography (CT)-guided high precision cone beam irradiation of animals carrying PARPi-resistant KB1P(M) tumours with low *Rev7* gene expression (Fig. 3c).

We then tested whether the processing of broken DNA ends requires ATM in *Rev7*-depleted cells, and found that inhibition of ATM using KU55933 efficiently suppresses DNA-damage-induced RAD51 foci and increases olaparib sensitivity (Extended Data Fig. 4h, i). Hence, the partial restoration of RAD51 focus formation in *Brca1*-deficient mammary tumour cells after DNA damage by inhibition of *Rev7* is ATM-dependent.

In contrast to the results with BRCA1-deficient cells, *Rev7* depletion in BRCA2-deficient cells did not result in PARPi resistance (Extended Data Fig. 5a–f). Furthermore, we did not observe increased PARPi

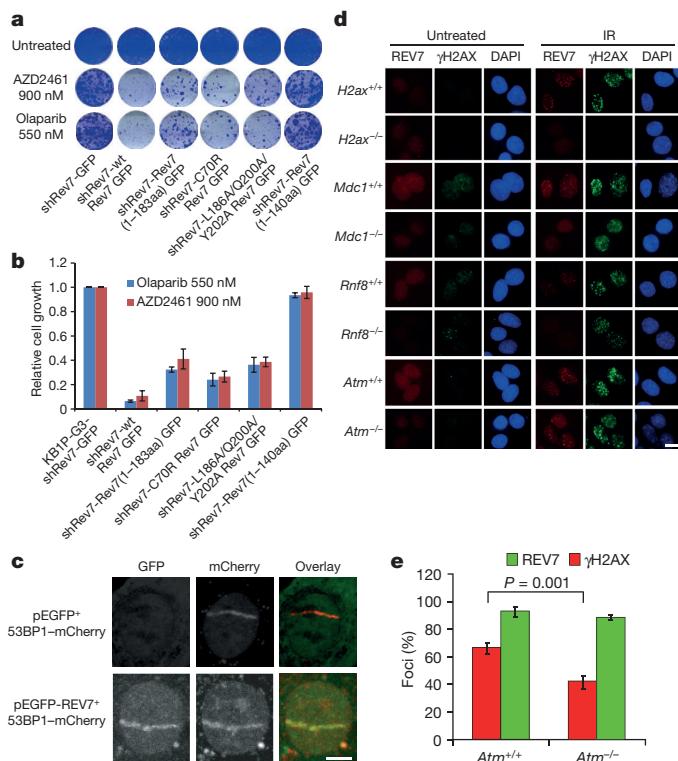


Figure 2 | Dissection of REV7 function and its dependent factors. **a, b**, Long-term clonogenic assay (**a**) and quantification (**b**) using KB1P-G3 cells transduced with the indicated constructs (wt *Rev7* stands for pMSCV-wt *Rev7*) and treatments. All groups were normalized to the absorbance of the shRev7-GFP control. The data represent the mean ± s.d. **c**, GFP-REV7 recruitment to sites of DNA damage (visualized by 53BP1-mCherry) was observed 5 min after 405 nm laser exposure (0.99 mW, 60% laser power, 50 s) in KB1P-B11 cells. pEGFP denotes a mammalian expression vector containing enhanced GFP. Scale bar, 5 μm. **d**, REV7 foci formation in *H2ax*^{-/-} (also known as *H2afx*^{-/-}), *Atm*^{-/-}, *Mdc1*^{-/-} and *Rnf8*^{-/-} mouse embryonic fibroblast (MEF) cells and their corresponding controls before and 4 h after 10 Gy ionizing radiation (IR). DAPI, 4',6-diamidino-2-phenylindole. Scale bar, 10 μm. **e**, Quantification of REV7 foci formation (>8 foci per cell) in *Atm*^{-/-} and *Atm*^{+/-} MEF cells. The quantification of foci-positive cells was performed by counting a total of 100 cells per sample. Data are presented as mean ± s.d. from three different experiments. *P* value calculated using the *t*-test.

resistance after *Rev7* inhibition in the BRCA1/2-proficient *p53*^{-/-} tumour cell line KP3.33 (Extended Data Fig. 5g–i). This indicates that REV7 works upstream of BRCA2 and is antagonized by BRCA1. We therefore tested whether DNA end resection is altered in the absence of *Rev7* in BRCA1-deficient cells. Accumulation of the single-strand binding protein, RPA, was used as a marker for the generation of single-stranded DNA (ssDNA). Cells were exposed to α particles¹², and BRCA1 deficiency resulted in a marked decrease in RPA-positive α tracks compared to BRCA1-proficient cells (Fig. 3d, e). REV7 depletion in the *Brca1*^{-/-} *p53*^{-/-} cells largely suppressed this defect in both KB1P-G3 and KB1P-B11 cells (Fig. 3e and Extended Data Fig. 6a). In addition to coating resected DNA ends, RPA also interacts with ssDNA gaps (for example, during replication)¹³. To exclude that the observed RPA accumulation reflected interaction with internal ssDNA gaps, we prevented DNA end resection (without influencing replication) by knocking down CTIP (ref. 14). This eliminated the increase in RPA-positive tracks induced by *Rev7* knockdown in *Brca1*^{-/-} *p53*^{-/-} cell lines (Fig. 3f and Extended Data Fig. 6b), without affecting cell cycle distribution (Extended Data Fig. 6c). We therefore conclude that increased resection and not binding to ssDNA gaps is responsible for RPA accumulation.

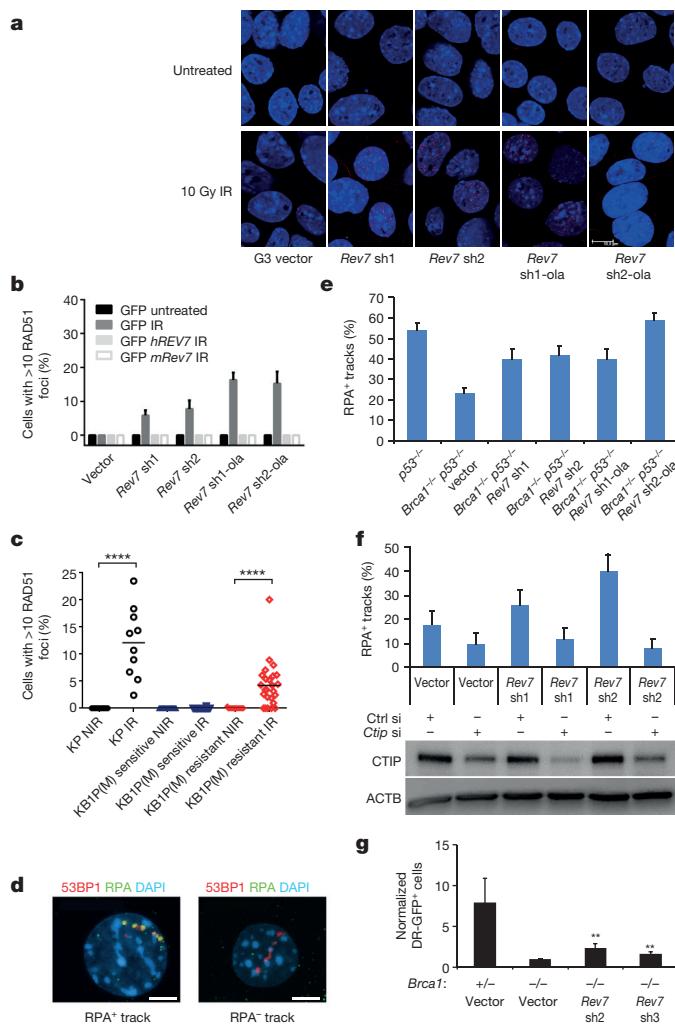


Figure 3 | The effect of REV7 inhibition on RAD51 and RPA focus formation of *Brca1*^{-/-} *p53*^{-/-} cells. **a**, RAD51 focus (red) formation in KB1P-G3 cells before and 5 h after 10 Gy ionizing radiation. Scale bar, 10 µm. **b**, Quantification of RAD51 foci in KB1P-G3 cells (with or without REV7 depletion) transfected with an empty vector (GFP) or vectors containing mouse or human *Rev7* or *REV7*, respectively. At least 150 GFP-positive cells were analysed per group in three independent experiments each. The data represent the mean ± s.d. IR denotes 5 h after 10 Gy ionizing radiation. **c**, *In situ* analysis of RAD51 foci in PARPi-resistant KB1P(M) tumours with low *Rev7* gene expression. KP (KP3.33) denotes mouse mammary tumour cell line (*p53*^{-/-}); IR denotes 2 h after 15 Gy ionizing radiation; NIR denotes no ionizing radiation. ***P < 0.0001, Mann–Whitney U test. **d**, **e**, Representative images of 53BP1-labelled α tracks in cells positive or negative for RPA (**d**) and quantification of RPA-positive tracks 2 h after ionizing radiation (**e**). KP (*p53*^{-/-}) or KB1P-G3 cells with or without *Rev7*-targeting shRNAs were tested. Scale bar, 5 µm. **f**, Quantification of RPA- and 53BP1-positive α tracks in KB1P-G3 cells transfected with non-targeting control (ctrl) short interfering RNAs (siRNAs) or siRNAs against mouse *Ctip*. CTIP protein expression of the indicated groups was checked by western blotting, with β-actin (ACTB) as a loading control. **g**, Quantification of HR using the DR-GFP reporter assay. GFP-positive cells normalized to the vector-transduced *Brca1*^{-/-} *p53* shRNA cells are shown. The data represent mean ± s.d. **P < 0.01, two-tailed t-test.

As *Rev7* loss could restore end resection in BRCA1-deficient cells, we analysed whether its depletion could restore full HR proficiency in this context. Using mouse embryonic stem (mES) cells with a *Brca1* selectable conditional knockout allele¹⁵, we observed that *Rev7* loss indeed prevented cell death of mES cells after BRCA1 deletion, and restored RAD51 focus formation upon DNA damage (Extended Data Fig. 6d–h). Moreover, we reproducibly observed a partial restoration of HR function

in the DR-GFP reporter assay for homologous recombination¹⁶ when *Rev7* was depleted in BRCA1-deficient mES cells (Fig. 3g).

Our data for REV7 are reminiscent of previous findings that 53BP1 loss occurs in subsets of human breast carcinomas¹⁵ and can also restore HR to BRCA1-deficient cells^{7,15,17}. As with 53BP1, we found a frequent aberrant reduction or loss of the REV7 protein in human triple-negative breast carcinomas (Extended Data Fig. 7). Addressing the relationship between REV7, 53BP1 and the 53BP1 non-homologous end-joining (NHEJ) effector protein, RIF1 (ref. 18), we found that *Rev7* deficiency did not compromise the formation of endogenous 53BP1 or RIF1 foci (Extended Data Fig. 8a–d). By contrast, endogenous REV7 foci or laser-induced stripes were absent in 53BP1-depleted mouse and human cells (Fig. 4a and Extended Data Fig. 4c, d), strongly suggesting that REV7 acts downstream of 53BP1. This is also consistent with our results that PARPi resistance is not increased when both *Rev7* and *53bp1* are depleted (data not shown). Despite such strong evidence for a cooperative role for REV7 and 53BP1, we did not detect REV7 in 53BP1 immunocomplexes isolated from untreated cell lysates, or ATM-phosphorylated 53BP1 immunocomplexes containing RIF1 that were induced by DNA damage¹⁹ (Extended Data Fig. 8e and data not shown). Although the intricacies of the interactions remain to be determined, REV7 recruitment to DNA damage sites by 53BP1 may result from indirect interactions or an activity elicited by 53BP1 protein complexes in chromatin at DSB sites.

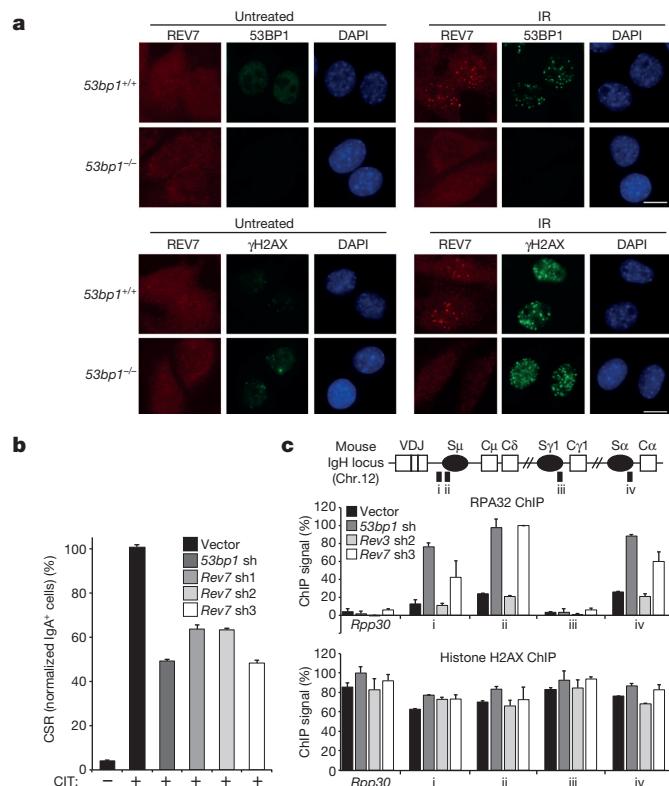


Figure 4 | REV7 is a downstream effector of 53BP1 on inhibiting end resection and promoting CSR. **a**, REV7 foci formation in 53bp1^{-/-} and 53bp1^{+/+} MEF cells before and 4 h after 10 Gy ionizing radiation. Scale bar, 10 µm. **b**, Quantification of CSR to IgA of shRNA-transduced CH12 cells 40 h after stimulation (CIT denotes CD40 antibody, IL-4 and TGF-β1). Data represent mean ± s.d. from two independent experiments performed in triplicate. **c**, Schematic of IgH locus shows relative positions of quantitative PCR amplicons used in ChIP experiments. A control non-IgH locus (*Rpp30*) was also examined. Indicated CH12 cell lines stimulated for 30 h with CIT were subjected to ChIP with IgG (control), histone H2AX and RPA32 monoclonal antisera. After background subtraction, values were normalized to the DNA input signals, followed by the maximum value in each data set. Mean signals, two replicates ± s.e.m.

To examine whether REV7, like 53BP1 (ref. 18), also promotes NHEJ of DSBs during class switch recombination (CSR), we depleted *Rev7* transcripts in the mouse CH12 B-cell line, which after stimulation undergoes CSR from IgM to IgA at a high rate²⁰. Efficient *Rev7* knock-down was achieved using several shRNAs, reducing CSR to levels comparable with 53BP1-depleted cells when compared to control-depleted cells (Fig. 4b and Extended Data Fig. 9a, b). Moreover, these defects were not accompanied by defects in cell proliferation, *Aid* (also known as *Aicda*), or germ-line transcript (μ GLT/ α GLT) expression (Extended Data Fig. 9c–e). Conditional REV3-ablation has been reported to reduce CSR efficiency in B lymphocytes²¹, suggesting that distinct from 53BP1, REV7 might participate in CSR through Pol ζ function. This considered, we reasoned that we might separate the function of REV7 from that of Pol ζ during CSR, at the level of DSB resection inhibition. To this end, RPA enrichment was measured by chromatin immunoprecipitation (ChIP) at control and several IgH loci in control, 53BP1-, REV3- or REV7-depleted CH12 lines stimulated to undergo CSR. Consistent with the role of 53BP1 in resection inhibition, 53BP1-depletion was accompanied by a 3–5-fold enrichment of RPA ChIP signal specifically in donor (S μ) and acceptor (S α) IgH switch (S) regions where DSBs occur during IgM to IgA CSR, but not at an IgH S γ 1 locus or a control non-IgH *Rpp30* locus (Fig. 4c). Notably, these defects were closely mimicked after REV7 depletion, whereas shRNA-mediated REV3 depletion yielded no detectable increase in RPA IgH S-region enrichment, despite diminishing CSR efficiency (Extended Data Fig. 9f) as expected²¹. Importantly, at each locus equivalent ChIP signals for total histone were obtained between cell lines. Thus, our data support a Pol ζ -independent function of REV7 in inhibiting the nucleolytic processing of DSBs generated during CSR.

Our data uncover a crucial role for REV7 in regulating DSB repair. REV7 depletion restores homology-directed DNA repair of BRCA1-deficient cells resulting in PARPi resistance. We attribute this result to the inhibitory effect of REV7 on DNA end resection. Like 53BP1 together with RIF1 and PTIP (ref. 18), REV7 may function as a NHEJ factor that performs a regulatory role in DSB repair pathway choice. This provides new insight into the versatile functions of REV7 in addition to its role in translesion synthesis^{9,22}, and helps to explain how HR can be partially restored in the absence of BRCA1.

Online Content Methods, along with any additional Extended Data display items and Source Data, are available in the online version of the paper; references unique to these sections appear only in the online paper.

Received 17 February 2014; accepted 13 February 2015.

Published online 23 March 2015.

- Roy, R., Chun, J. & Powell, S. N. BRCA1 and BRCA2: different roles in a common pathway of genome protection. *Nature Rev. Cancer* **12**, 68–78 (2012).
- Lee, J.-M., Ledermann, J. A. & Kohn, E. C. PARP Inhibitors for BRCA1/2 mutation-associated and BRCA-like malignancies. *Ann. Oncol.* **25**, 32–40 (2013).
- Farmer, H. *et al.* Targeting the DNA repair defect in BRCA mutant cells as a therapeutic strategy. *Nature* **434**, 917–921 (2005).
- Lord, C. J. & Ashworth, A. Mechanisms of resistance to therapies targeting BRCA-mutant cancers. *Nature Med.* **19**, 1381–1388 (2013).
- Ang, J. E. *et al.* Efficacy of chemotherapy in BRCA1/2 mutation carrier ovarian cancer in the setting of PARP inhibitor resistance: a multi-institutional study. *Clin. Cancer Res.* **19**, 5485–5493 (2013).
- Chun, A. C.-S., Kok, K.-H. & Jin, D.-Y. REV7 is required for anaphase-promoting complex-dependent ubiquitination and degradation of translesion DNA polymerase REV1. *Cell Cycle* **12**, 365–378 (2013).
- Jaspers, J. E. *et al.* Loss of 53BP1 causes PARP inhibitor resistance in *Brca1*-mutated mouse mammary tumors. *Cancer Discov* **3**, 68–81 (2013).

- Listovsky, T. & Sale, J. E. Sequestration of CDH1 by MAD2L2 prevents premature APC/C activation prior to anaphase onset. *J. Cell Biol.* **203**, 87–100 (2013).
- Gan, G. N., Wittschieben, J. P., Wittschieben, B. Ø. & Wood, R. D. DNA polymerase zeta (pol zeta) in higher eukaryotes. *Cell Res.* **18**, 174–183 (2008).
- Hara, K. *et al.* Crystal structure of human REV7 in complex with a human REV3 fragment and structural implication of the interaction between DNA polymerase zeta and REV1. *J. Biol. Chem.* **285**, 12299–12307 (2010).
- Khalaj, M. *et al.* A missense mutation in Rev7 disrupts formation of Pol ζ , impairing mouse development and repair of genotoxic agent-induced DNA lesions. *J. Biol. Chem.* **289**, 3811–3824 (2014).
- Stap, J. *et al.* Induction of linear tracks of DNA double-strand breaks by alpha-particle irradiation of cells. *Nature Methods* **5**, 261–266 (2008).
- Fanning, E., Klimovich, V. & Nager, A. R. A dynamic model for replication protein A (RPA) function in DNA processing pathways. *Nucleic Acids Res.* **34**, 4126–4137 (2006).
- Sartori, A. A. *et al.* Human CtIP promotes DNA end resection. *Nature* **450**, 509–514 (2007).
- Bouwman, P. *et al.* 53BP1 loss rescues BRCA1 deficiency and is associated with triple-negative and BRCA-mutated breast cancers. *Nature Struct. Mol. Biol.* **17**, 688–695 (2010).
- Pierce, A. J., Johnson, R. D., Thompson, L. H. & Jasin, M. XRCC3 promotes homology-directed repair of DNA damage in mammalian cells. *Genes Dev.* **13**, 2633–2638 (1999).
- Bunting, S. F. *et al.* 53BP1 inhibits homologous recombination in *Brca1*-deficient cells by blocking resection of DNA breaks. *Cell* **141**, 243–254 (2010).
- Zimmermann, M. & de Lange, T. 53BP1: pro choice in DNA repair. *Trends Cell Biol.* **24**, 108–117 (2014).
- Chapman, J. R. *et al.* RIF1 is essential for 53BP1-dependent nonhomologous end joining and suppression of DNA double-strand break resection. *Mol. Cell* **49**, 858–871 (2013).
- Muramatsu, M. *et al.* Class switch recombination and hypermutation require activation-induced cytidine deaminase (AID), a potential RNA editing enzyme. *Cell* **102**, 553–563 (2000).
- Schenten, D. *et al.* Pol ζ ablation in B cells impairs the germinal center reaction, class switch recombination, DNA break repair, and genome stability. *J. Exp. Med.* **206**, 477–490 (2009).
- Kikuchi, S., Hara, K., Shimizu, T., Sato, M. & Hashimoto, H. Structural basis of recruitment of DNA polymerase ζ by interaction between REV1 and REV7 proteins. *J. Biol. Chem.* **287**, 33847–33852 (2012).

Supplementary Information is available in the online version of the paper.

Acknowledgements We thank B. Gerritsen, P. Halonen, B. Morris, T. Halazonetis and O. Kallioniemi for advice on the DDR shRNA library, A. Gasparini and G. Borst for their assistance with the cone beam micro-irradiator, R. Kanaar for his RAD51 antibody, J. Jacobs for the pMSV-GFP vector, and M. O'Connor for olaparib and AZD2461. This work was supported by the Netherlands Organization for Scientific Research (NWO-Toptalent to J.E.J. and NWO-VIDI to S.R.), the Dutch Cancer Society, CTMM Breast Care, the Swiss National Science Foundation, and the European Union (EU) FP6 Integrated Project CHEMORES and FP7 Project DDResponse. Work in J.R.C.'s group is funded by the Wellcome Trust. The work in the J.B.'s laboratories was funded by the Danish Cancer Society, the Danish Council for Independent Research, the Lundbeck Foundation and the Czech National Program of Sustainability. S.J.B. is funded by Cancer Research UK and an ERC Advanced Investigator Grant (RecMitMei) and is a Royal Society Wolfson Research Merit Award Holder.

Author Contributions G.X. and S.R. designed the study, performed experiments and wrote the manuscript; I.B. and D.C.v.G. designed and performed the RPA foci analysis; J.R.C., A.S. and S.J.B. performed and planned the CSR assay and RIF1-associated data; J.C. and J.Y. designed and performed the experiments using MEFs; J.E.J., A.K. and W.S. assisted with the mouse intervention studies; E.G. established the *in vivo* RAD51 analysis that K.J. and B.v.d.B. quantified; P.H.N.C. designed and M.B. helped in generating the REV7 mutants; Pe.B., M.P. and J.J. helped in designing the shRNA screen and performed experiments using mES cells; M.M. performed the laser stripe assays, and K.W. performed co-immunoprecipitations. D.W. helped to visualize GFP-REV7 recruitment; M.N., I.d.R. and J.d.R. carried out the RNA sequencing (RNA-seq) data analysis, Jirina B. established, carried out and evaluated the REV7 immunohistochemistry, P.C.S. helped with the analysis of the immunohistochemistry data, Jiri B. and Pi.B. advised on experiments and manuscript revisions, and all authors discussed and approved the manuscript.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to S.R. (sven.rottenberg@vetsuisse.unibe.ch).

METHODS

Cell culture and reagents. KB1P-B11 (B11) and KB1P-G3 (G3) cell lines were derived from a *Brcal*^{-/-} *p53*^{-/-} mouse mammary tumour as described⁷. KB2P-1.21 and KB2P-3.4 cell lines originate from a *Brcal*^{-/-} *p53*^{-/-} mouse mammary tumour, and KP3.33 cell line from a *p53*^{-/-} mouse mammary²³. These cell lines were cultured in DMEM/F-12 (Life Technologies) supplemented with 10% FCS, 50 U ml⁻¹ penicillin, 50 ng ml⁻¹ streptomycin, 5 µg ml⁻¹ insulin (Sigma), 5 ng ml⁻¹ epidermal growth factor (Life Technologies) and 5 ng ml⁻¹ cholera toxin (Gentaur) under low oxygen conditions (3% O₂, 5% CO₂, 37 °C). SUM149PT cells were grown in RPMI1640 (Life Technologies) supplemented with 10% FCS, under normal oxygen conditions (21% O₂, 5% CO₂, 37 °C). U2OS, phoenix, 293T cells were cultured in DMEM (Life Technologies) supplemented with 10% FCS, under normal oxygen conditions (21% O₂, 5% CO₂, 37 °C). mES cells with a selectable conditional BRCA1 deletion (*R26*^{CreERT2/wt}; *Brcal*^{SCo/Δ})¹⁵ were cultured on gelatin-coated plates in 60% buffalo red liver cell-conditioned medium supplied with 10% FCS, 0.1 mM β-mercaptoethanol (Merck) and 1 × 10³ U ml⁻¹ ESGRO LIF (Millipore) under normal oxygen conditions (21% O₂, 5% CO₂, 37 °C). CH12 cells (CH12F3-2) were cultured in RPMI-1640 medium (Sigma) supplemented with 10% FBS, 5% NCTC 109 medium (Sigma), 50 µM β-mercaptoethanol, 50 U ml⁻¹ penicillin and 50 ng ml⁻¹ streptomycin under normal oxygen conditions. Olaparib and AZD2461 were provided by AstraZeneca; KU-55933 (KuDOS) was bought from Selleckchem (S1092).

Lentivirus-based transduction of cells with shRNA. Glycerol stocks of shRNA hairpins were obtained from the Sigma Mission library (TRC 1.0) and isolation of plasmids was carried out with the high pure plasmid Mini Kit or Genopure maxi kit (Roche). 293T cells were seeded 16 h before transfection. For each 10-cm dish, 0.5 ml 2× HBS (8.18 g NaCl, 0.2 g Na₂HPO₄·7H₂O, 5.95 g HEPES in 500 ml MilliQ water at pH 7.01) was added into a sterile falcon tube. In another sterile falcon tube, 6 µg plasmid DNA of interest, 2 µg pHCMV-G envelope vector (pMD.G), 2 µg pRSV-Rev, 2 µg packaging vector pMDl/g/pRRE, 250 µl 0.5 M CaCl₂ and distilled water were added to bring up to 0.5 ml. The CaCl₂/plasmid DNA mix was added to the 2× HBS and incubated for 20 min and then added to the cells. Medium was refreshed after 6 h and another 18 h, respectively. The supernatant of 293T cells containing lentivirus was collected after 24 h to infect cells with polybrene (6 µg ml⁻¹) for 12 h. The medium was refreshed after lentivirus infection and the cells were selected with puromycin.

Individual shRNA vectors used were collected from the TRC library. Mouse *Rev7*: sh1: TRCN000012844_CCAGTGGAGAAGTTGTCTTT; sh2: TRCN 0000012846_CATCTTCCAGAACGCGAACAGAA; sh3: TRCN0000012847_GATACAGGTATCAAGGACTT; human *REV7*: sh1: TRCN000006569_CCCCTGA TTCCAAGTGTCTTA; sh2: TRCN000006570_CCCGGAGCTGAATCAGTA TAT; sh3: TRCN000006571_CCCAGTGGAGAAATTCTGTCTT; sh4: TRCN00 0006573_CATCTTCCAGAACGCAAGAA; mouse *53bp1*: (puromycin) sh: TRCN0000081778_GCTATTGTGGAGATTGTGTTT; (neomycin) sh: same sequences as above; human *53BP1*: sh1: TRCN0000018866_CCAGTGTGATTAGT ATTGATT; sh2: TRCN0000018865_GATACTTGGCTTACTGGTTT; mouse *p53*: (neomycin) TRCN0000054551_AGAGTATTTCACCCCTCAAGAT; mouse *Rev1*: sh1: TRCN0000120298_GCCGAGATCAACTATGAAATA; sh2: TRCN 0000120297_CAGCAGTGTGTTGAGGTATT; mouse *Rev3*: sh1: TRCN0000 119969_CCGTCACATTAGTGAGACTAT; sh2: TRCN0000119970_GCCCAC ATACACTTCTTCTT.

Loss-of-function screen. In total, 1,976 lentiviral hairpins (pLKO.1) from the Sigma Mission library (TRC Mm 1.0) that target 391 mouse genes involved in the DNA damage response were selected (see Supplementary Table 1). This library was used to generate pools of lentiviral shRNA in 293T cells to infect target cells. After infection, the cells stably expressing integrated shRNA were selected with puromycin. Cells with HR restoration were selected with a high concentration of olaparib (500 nM, about 100-fold the IC₅₀), which killed cells of the empty vector group. Surviving cells were pooled and genomic DNA was extracted using the Gentra Puregene kit according to the manufacturer's protocol (Qiagen). shRNA inserts were retrieved from 50 ng genomic DNA by PCR amplification (PCR1 and PCR2) using the following conditions: (1) 95 °C, 5 min; (2) 95 °C, 30 s; (3) 60 °C, 30 s; (4) 72 °C, 1 min; (5) go to step (2), 20 cycles; (6) 72 °C, 5 min; (7) 4 °C. The PCR reaction system were as follows: 0.6 µl DMSO, 4.0 µl phusion HF buffer 5×, 0.4 µl dNTPs, 1.0 µl primer f (10 mM), 1.0 µl primer r (10 mM), 11.8 µl mQ, 0.2 µl phusion, 1.0 µl gDNA for PCR1 or 1 µl PCR1 products for PCR2. Adaptors and indexes for deep sequencing (Illumina HiSeq 2000) were incorporated into PCR primers as follows: PCR1 forward: PCR1_01_PLKO1_f_Integration determination_1, 5'-AC ACTCTTCCCTACACGACGCTCTCCGATCTGTAGCCCTGTGGAAAAG GACGAAACACCGG-3'; PCR1_02_PLKO1_f_Integration determination_2, 5'-ACACTCTTCCCTACACGACGCTCTCCGATCTGTAGCCCTGTGGAAAAG GACGAAACACCGG-3'; PCR1_03_PLKO1_f_untreated_1, 5'-ACACTCTT CCCTACACGACGCTCTCCGATCTACTGTCTGTGGAAAGGACGAA ACACCGG-3'; PCR1_04_PLKO1_f_untreated_2, 5'-ACACTCTTCCCTAC

ACGACGCTCTCCGATCTATTGGCCTGTGGAAAGGACGAAACACCGG -3'; PCR1_05_PLKO1_f_olaparib_1, 5'-ACACTCTTCCCTACACGACGCTC TTCCGATCTGATCTGCTGTGGAAAGGACGAAACACCGG-3'; PCR1_06_PLKO1_f_olaparib_2, 5'-ACACTCTTCCCTACACGACGCTCTCCGATCTGATCTGATCTGCTGTGGAAAGGACGAAACACCGG-3'; PCR1_07_PLKO1_f_olaparib_3, 5'-ACACTCTTCCCTACACGACGCTCTCCGATCTGATCTGATCTGCTGTGGAAAGGACGAAACACCGG-3'; PCR1_08_PLKO1_f_olaparib_4, 5'-ACACTCTTCCCTACACGACGCTCTCCGATCTAAGCTACTGTGTGGAA AGGACGAAACACCGG-3'; PCR1 reverse: P7_pLKO1_r, 5'-CAAGCAGAAC CGGCATACGGAGATTTCTTCCCTGCACTGTACCC-3'; PCR2 forward: P5_IlluSeq, 5'-AATGATACGGCACCACGGAGATCTACACTCTTCCCTA CACGACGCTCTCCGATCT-3'; PCR2 reverse: P7, 5'-CAAGCAGAACACGG CATAACGGAGAT-3'.

PCR2 products were purified using the PCR purification kit from Qiagen. The shRNA stem sequence was segregated and aligned to the TRC library. The reads of different hairpins were counted and the following criteria were used to select the top hits for further validation: (1) hairpins targeting the same gene in survival clones should have at least 1 × 10⁴ reads (total ~6 × 10⁶ reads); (2) at least two different hairpins targeting the same gene should be present; (3) hairpins in resistant clones should be highly enriched (>8-fold) in cells after olaparib selection; and (4) hairpins should be present in 4 out of 4 independent screens.

PARPi treatment study. Long-term clonogenic assay: on day 0, 1.5 × 10⁴ (B11) or 1 × 10⁴ (G3, KB2P-1.21 or KB2P-3.4) or 6 × 10³ (KP3.33) cells were seeded per well with PARPi (or untreated control) into six-well plates. The medium of the PARPi treatment groups was refreshed with PARPi every 4 days. On day 5, the untreated control group was stopped and the PARPi treatment groups were stopped after another 2–3 weeks and stained with 0.1% crystal violet. Using the SUM149PT cells, 4 × 10⁴ cells were seeded per well with olaparib (or untreated control) into 12-well plates. The medium of the PARPi treatment groups was refreshed with olaparib every 4 days. On day 6, the untreated control group was stopped and the olaparib treatment groups were stopped on day 8 and stained with 0.1% crystal violet. Quantification of the clonogenic assay was done by determining the absorbance of crystal violet at 590 nm.

Short-term clonogenic assay: on day 0, 4 × 10² (G3) cells were seeded per well with olaparib (or untreated control) into 6-well plates. On day 4, the medium was refreshed with olaparib or untreated control. On day 8, all the groups were stained with Leishman dye and quantification was done by the relative colony numbers.

ATM and PARP inhibitor combination study. On day 0, 1 × 10⁴ (G3) cells were seeded per well into 6-well plates and then ATM inhibitor or olaparib or their combination was added. The medium was refreshed every 3 days with the different drugs. For the combination therapy groups, ATM inhibitor was applied for 6 days. On day 5, ATM inhibitor alone and untreated control groups were stopped and the other groups were stopped on day 12 and stained with 0.1% crystal violet.

Constructs. Human *REV7* was amplified by PCR using Platinum Taq polymerase (Invitrogen) from U2OS cDNA using the following primers 5'-ATAGAACATC AATGACCACGCTCACACGACAAGAC -3' and 5'-ATATGGTACCATCAGC TGCCTTATGAGCCCGC-3'. Mouse *Rev7* was amplified from mouse lung cDNA in two parts to introduce silent mutations, making it resistant to *Rev7* shRNA2 (m*Rev7R*). The following primers were used for part A: [IB11m] 5'-ATATGAAT TCGATGACCACCCCTCACGCGC-3' [IB14m] 5'-TACTTCTCCGTTCTGA AAGATGCCAACCGGGTA-3' and part B: [IB12m] 5'-ATATGGTACC-AT-G CTGTTCTATGCCTCGCT-3' [IB13m] 5'-GGGCATTTCAGAAACCGG AAGAAGTACAACGTGC-3'. Equal parts of both PCR reactions were mixed and used for a PCR using IB11m and IB12m to create the complete mouse *Rev7* sequence including the silent mutations.

The PCR product and pEGFP-C1 vector were digested using EcoRI and KpnI and ligated using the T4 DNA Ligase (Roche) to generate pEGFP-h*REV7* and pEGFP-m*REV7R*.

Using pEGFP-m*REV7R* as a template, pEGFP-C1-based *REV7* truncated constructs was amplified by PCR using the following primers: forward: 5'-ATATGAATTCGATGACCACCCCTCACGCGC-3', reverse: m*REV7R* (1–55aa): 5'-ATATGGTACCGGACATCTGAACCGGGAC-3'; m*REV7R* (1–81aa): 5'-ATATGGTACCCCTCACATCGTTCTCCAGG-3'; m*REV7R* (1–110aa): 5'-ATATGGTACCGATGGACAGCAAGGGAGGC-3'; m*REV7R* (1–140aa): 5'-ATATGGTACCGITGTGATCCAGGACAGC-3'; m*REV7R* (1–183aa): 5'-ATATGGTACCGTGTGATCCAGGACATCTG-3'. pEGFP-REV7 mutants (shRNA-resistant) were ordered as gBlocks (IDT) and cloned into the pEGFP-C1 vector using EcoRI and KpnI restriction enzymes and quick ligase (NEB). The DNA sequences that were ordered as gBlocks are: m*REV7_shRNAresistant* (m*REV7R*): CGCCCGA AATTCCGCCACCATGACCACCCCTACGCGCCAAGACCTCAACTTGGC CAAGTGGTGGCTGACGTGCTCTCCGAGITCTCTGGAGGTGGCCGTGCAC CTGATTCTCTATGTGCGCGAGGTCTACCCGGTGGGCATCTTCAGAAA CGGAAGAAGTACAACGTGCCGGTTCAGATGTCTGTCAACCGGAGCTG

AACCAAGTACATCCAGGACACACTCCACTGCGTCAAACCTCTCCTGGAG
AAGAACGATGTGGAGAAGGTGGTGGTGGATTTGGATAAGGAACAC
CGCCCAGTGGAGAAGTTGTCTTGAGATCACTCAGCCTCCCTGCTG
TCCATCAATTCAAGACTCCCTCGTCTCATGTGGAGCAGCTGCTCGAG
CCTTCATCCTTAAGATTAGTGTGTGATGCTGCTGGATCACAAACC
CTCCAGGCTGCACATTACAGTCCTCGACACAAGAGAACAGCTGCTA
CTCGAAACATGGAGAAGATAACAGGTACATCAAGGACTTCCCATGGATC
CTGGCAGATGAACAGGATGTCACATGACGACCCCCGCTTGATACCC
CTAAAAACCATGACGTCGGACATTAAAGATCAGCTACAGTTGAA
GAGCGAGGCCATAAGAACACAGCTGAGGTACCCGGG; mREV7_shRNA
resistant_L186A/Q200A/Y202A: CGCCCGAATTCGCCACCATGACC
CTCACGCCAAGACCTCAACTTGGCCAAGTGGCTGAGCTGCT
TCCGAGTCTGGAGGTGGCGTGCACCTGATTCTATGTGCGGAG
GTCTACCCGGTGGCATCTTCAGAAACGGAGAAGTACAACGTGCC
GGITCAGATGTCTGTCAACCGGAGCTGAACCAAGTACATCCAGGACAC
ACTCCACTGCGTCAAACCTCTCCTGGAGAAGAACGATGTGGAGAAGG
TGGTGGTGGTATTTGGATAAGGAACACCGGCCAGTGGAGAAGT
GTCTTGAGATCACTCAGCCTCCCTGCTGTCATCAATTCAAGACTCCC
TCCGTCTCATGTGGAGCAGCTGCTGAGCCTCATCCTTAAGATTAG
TGTGTGTGATGCTGCTGGATCACAACCCCTCAGGCTGCACATTAC
AGTCCTCGACACAAGAGAAGCTGACTCGAAACATGGAGAAGA
TACAGGTATCAAGGACTTCCATGGATCCTGGCAGATGAACAGGATG
TCCACATGACGACCCCCCGCGTACACCCCTAAAACCATGACGTCGG
ACATTAAAGATGGCTCGTGTGAAGAGCGAGCGATAAGAAC
AGCTGAGGTACCCGGG; mREV7_shRNAresistant_C70R: CGCCCGA
TTCGCCACCATGACCAACCTCACGGCCAAGACCTCAACTTGGCA
AGTGGTGGCTGACGTGCTCCGAGTCTGGAGGTGGCGCTGCACCT
GATTCTCATGTGCGCAGGTCTACCCGGTGGCATCTTCAGAAACG
GAAGAAGTACAACGTGCCGGTTCAGATGTCTGTCAACCGGAGCTGAA
CCAGTACATCCAGGACACACTCCACCGCGTCAAACCTCTCCTGGAGAA
GAACGATGTGGAGAAGGTGGTGGTGGTGGATTTGGATAAGGAACACC
GCCCAAGTGGAGAAGTTGTCTTGAGATCACTCAGCCTCCCTGCTG
CCATCAATTCAAGACTCCCTCTGTCTCATGTGGAGCAGCTGCTCGAG
CCTTCATCCTTAAGATTAGTGTGTGATGCTGCTGGATCACAAACC
CTCCAGGCTGCACATTACAGTCCTCGACACAAGAGAACGCTGCTA
CTCGAAACATGGAGAAGATAACAGGTATCAAGGACTTCCATGGATC
CTGGCAGATGAACAGGATGTCACATGACGACCCCCGGTGTGATACCC
CTAAAAACCATGACGTCGGACATTAAAGATGTCAGCTACGGTGA
GAGCGAGGCCATAAGAACAGCTGAGGTACCCGGG. Using pEGFP-
mREV7R, C70R and L186A/Q200A/Y202A mutants as templates, Gateway compatible pMSCV-GFP or pLenti6-UBC (Invitrogen)-based REV7 truncated or REV7 mutated constructs were amplified by PCR using the following primers: forward: 5'-GGGGACAACCTTGTACAAAAAGTTGGCATGACCACCTCACCGCG
CAA-3'; reverse: mREV7R (full-length for wt REV7, C70R REV7, L186A/Q200A/Y202A REV7): 5'-GGGGACAACCTTGTACAAGAAAGTTGGTATTCA
TGTTCTTATGCGCTCGCTC-3'; mREV7R(1-140aa): 5'-GGGGACAACCTTGTACAAGAAAGTTGGTATTCA
GTACAAGAAAGTTGGTATTCAAGGACAGC-3'; mREV7R
(1-183aa): 5'-GGGGACAACCTTGTACAAGAAAGTTGGTATTCA
GCATGTGGACATCTG-3'. PCR products were purified using PCR purification kit (Roche) and then subjected to BP (BP-clonase II, Invitrogen) and LR (LR-clonase II, Invitrogen) reaction according to manufacturer's instructions. Wild-type 53BP1 and 53BP1^{20AQ} constructs were previously described¹⁹. All the constructs were verified by sequencing.

siRNA, cDNA transfection and cDNA transduction. SMARTpool siRNAs targeting mouse *Ctip* (siGENOME: M-055713-02, Thermo Scientific) and non-targeting control were transfected into cells using Dharmacon 1 transfection reagent (Thermo Scientific). After 48 h, cells were subjected to western blot and α track assay. siRNAs against 53bp1, *Rnf8* or *Rnf168* were transfected into cells using Lipofectamine RNAiMax according to the manufacturer's instruction. The sequences of 53bp1 siRNA: 5'-GAGAGCAGAUAGCUUUUA-3' (ref. 24); *Rnf8* siRNA: 5'-GGAC
AAUUAUGGACAACAATT-3' (ref. 25); *Rnf168* siRNA: 5'-GGCGAAGAGCG
AUGGAGGATT-3' (ref. 26); GFP siRNA: 5'-GGCUACGUCCAGGAGCGCACC
TT-3'. Immunofluorescence and western blotting analysis was done 64 h after transfection.

For pEGFP-based constructs, transient transfection was done using X-treme GENE HP DNA transfection reagent (Roche) according to the manufacturer's manual. GFP-positive cells were sorted by flow cytometry and subjected to western blotting and immunofluorescence staining.

For pMSCV-GFP (retrovirus)-based constructs, transient transfection was done using X-treme GENE HP DNA transfection reagent in phoenix cells and the medium was refreshed after 24 h. Retroviruses were collected 48 h after transfection, and then target cells were infected for two consecutive periods of 12 h using fresh

virus. The medium was refreshed after retrovirus infection and the cells were selected with blasticidin.

For pLenti6-UBC (lentivirus) based constructs, together with the packaging plasmids p59, p60 and p61, transient transfection was done using X-treme GENE HP DNA transfection reagent in 293T cells and the medium was refreshed after 24 h. Lentiviruses were collected 48 h after transfection and then target cells were infected for two consecutive periods of 12 h using fresh virus. The medium was refreshed after lentivirus infection and the cells were selected with blasticidin.

Mice, generation of PARPi-resistant mouse mammary tumours. All mouse experiments were approved by the Animal Experiments Review Board of the Netherlands Cancer Institute (Amsterdam), complying with Dutch legislation. Olaparib-resistant KB1P(M)- and AZD2461-resistant KB1P mouse mammary tumours were generated as described⁷. In this study we analysed a total of 55 PARPi-resistant and 52 PARPi-sensitive tumours derived from 13 individual KB1P(M) and 10 individual KB1P donor tumours.

To generate mouse mammary tumours from cell lines, 5×10^5 cells were orthotopically transplanted into 6-week-old female wild-type FVB/N_Ola129 mice as reported previously⁷. Mice were randomized to the PARPi or untreated control groups.

RT-qPCR. Total RNA was isolated from cells using the high pure RNA isolation kit (Roche). cDNA was made from 1 μ g RNA with the GoScript Reverse Transcription system (Promega). For the quantitative PCR cDNA, Primer Fw&Rv (400 nM) and Lightcycler 480 SYBR Green I Master (Roche) were applied in a Lightcycler 480 96-well plate (Roche). The SYBR green signals were measured with Lightcycler 480 II (Roche). The C_p value of the gene of interest was subtracted from the house-keeping gene. This value was put in the power of 2 and this was also done for the s.d. The primer sequences used in this study are as follows: Mouse *Hprt* forward: 5'-CTGGTAAAAGGACTCTCG-3'; mouse *Hprt* reverse: 5'-TGAAGTACTCA TTATAGTCAAGGGCA-3'; mouse *Rev7* forward: 5'-ACACTCCACTGCGTCA AAC-3'; mouse *Rev7* reverse: 5'-AAAGACAAACTCTCCACTGGC-3'; mouse *Rev1* forward: 5'-ACAGGATTGCTGGTGCCTGT-3'; mouse *Rev1* reverse: 5'-TGAAGTCCCGCTTCTCTC-3'; mouse *Rev3* forward: 5'-AAGAGAT GTCACAGACAGGCC-3'; mouse *Rev3* reverse: 5'-AGTAGACAGCGCTG TTGTGC-3'; mouse α GLT forward: 5'-GACATGATCACAGGCACAGG-3'; mouse α GLT reverse: 5'-TTCCCCAGGTACATTCATCGT-3'; mouse μ GLT forward: 5'-TAGTAAGCGAGGCTCTAAAAGCAT-3'; mouse μ GLT reverse: 5'-AGAACAGTCCAGTGTAGGCAGTAGA-3'; mouse *Aid* forward: 5'-GAAA GTCACCGCTGGAGACCG-3'; mouse *Aid* reverse: 5'-TCTCATGCCGCTCCCT GG-3'. Human *HPRT*-P1 (primers pair 1) forward: 5'-GCAGACTTGTCTTC CTTGG-3'; human *HPRT*-P2 reverse: 5'-ACACTCGTGGGGTCTTT-3'; human *HPRT*-P2 forward: 5'-TGCTCGAGATGTGATGAAGG-3'; human *HPRT*-P2 reverse: 5'-ATCCAGCAGGTCAAGAAG-3'; human *REV7*-P1 forward: 5'-TGCTGTCCATCAGCTCAGAC-3'; human *REV7*-P1 reverse: 5'-TCT TCTCCATGTGCGAGTG-3'; human *REV7*-P2 forward: 5'-GTCACACGAC AAGACCTCA-3'; human *REV7*-P2 reverse: 5'-GACCGGCACGTTGTACTT CT-3'; mouse 53bp1 forward: 5'-TCAGCCAAACAGGACAAGCA-3'; mouse 53bp1 reverse: 5'-GCAGAACTTCAGCAGCAAGG-3'.

Western blotting. Cells were washed with ice-cold PBS and lysed on ice for 30 min with RIPA lysis buffer supplemented with three protease inhibitors (P8340, P5726, P0044; Sigma-Aldrich). Protein concentration was determined by the Bradford Protein Assay Kit (Bio-Rad) and a calibration standard curve created from the BSA. The samples were prepared for loading by adding 4 \times sample buffer (Invitrogen) and heating the samples at 70 °C for 10 min. Total proteins were separated by SDS-PAGE on 3–8% Tris-acetate (for 53BP1, RIF1) or 4–12% Bis-Tris gradient gels (all others). Next, proteins in the gel were electrophoretically transferred to PVDF membrane (Millipore) (for REV7, ACTB, α -tubulin) or to NC membranes (Invitrogen) (all others) and then the membrane was blocked in 5% milk with Tris-buffered saline Triton X-100 buffer (100 mM Tris, pH 7.4, 500 mM NaCl, 0.1% Triton X-100) (TBS-T0.1%). Membranes were incubated with primary antibodies in 5% milk in TBS-T0.1% overnight at 4 °C. Horseradish peroxidase (HRP)-conjugated secondary antibody incubation was performed for 1 h at room temperature in 5% milk in TBS-T0.1% and signals were visualized by ECL. Primary antibodies used in this study were as follows: mouse anti-REV7 (612266, BD Biosciences), 1:5,000 dilution; rabbit anti-53BP1 (NB100-304, Novus), 1:1,000 dilution; rabbit anti-53BP1 (A300-272A, Bethyl), 1:5,000 dilution; rabbit anti-CTIP (ab70163, Abcam), 1:1,000 dilution; mouse anti- α -tubulin (T6074, Sigma), 1:5,000 dilution; mouse anti-ACTB (MAB1501R, Millipore), 1:5,000 dilution; rabbit anti-mouse RIF1 (SK1316) (ref. 19), 1:2,000 dilution; mouse anti-RNF8 (B-2, Santa Cruz), 1:1,000 dilution; rabbit anti-RNF168 (ref. 27), 1:5,000 dilution; mouse anti-GAPDH (1D4, GeneTex) 1:1,000 dilution. Secondary antibodies used in this study were as follows: polyclonal rabbit anti-mouse immunoglobulins/HRP (P0161, Dako), 1:10,000 dilution; polyclonal swine anti-rabbit immunoglobulins/HRP (P0217, Dako), 1:10,000 dilution.

Immunofluorescence. Cells were grown on glass coverslips (12 mm) in 24-well plates. To induce ionizing radiation-induced foci, cells were γ -irradiated (10 Gy) and compared to non-irradiated controls 5 h after ionizing radiation. For this purpose the cells were pre-extracted using cold CSK buffer (10 mM HEPES-KOH, pH 7.9, 100 mM NaCl, 300 mM sucrose, 3 mM MgCl₂, 1 mM EGTA and 0.5% (v/v) Triton X-100) on ice for 5 min and then cold CSS buffer (10 mM Tris, pH 7.4, 10 mM NaCl, 3 mM MgCl₂, 1% (v/v) Tween and 0.5% (w/v) sodium deoxycholate) on ice for 5 min. Cells were washed with PBS++ (PBS with 1 mM CaCl₂ and 0.5 mM MgCl₂) and fixed using 2% PFA/PBS++ for 20 min at room temperature. Fixed cells were washed three times with PBS++ and stored at 4 °C. The cells were incubated 20 min in 0.2% Triton X-100/PBS++ to be permeabilized. Then the cells were washed three times in staining buffer (PBS++: 1% BSA, 0.15% glycine, 0.1% Triton X-100), incubated for 30 min in staining buffer at room temperature, incubated with the first antibody for 2 h at room temperature in staining buffer, washed three times in staining buffer, incubated with the second antibody for 1 h at room temperature in staining buffer and washed three times in staining buffer. Next, the cells were counter-stained with DAPI for 5 min, washed in staining buffer, washed in PBS++, mounted in Vectashield and sealed with nail polish. Primary antibodies used in this study were as follows: rabbit anti-RAD51 (70-001, BioAcademia), 1:20,000 dilution; rabbit anti-53BP1 (A300-272A, Bethyl), 1:4,000 dilution. Secondary antibodies used in this study were as follows: Alexa Fluor 568 F(ab')₂ Fragment goat anti-rabbit (A21069, Invitrogen), 1:1,000 dilution; Alexa Fluor 488 anti-mouse antibody (A11001, Invitrogen), 1:1,000 dilution.

For REV7 staining in MEF cells, cells cultured on coverslips were treated with 10 Gy IR and allowed to recover for 4 h. Cells were then washed with PBS, pre-extracted with 0.5% Triton X-100 solution for 3 min and fixed with 3% paraformaldehyde for 12 min. Coverslips were washed with PBS and then immunostained with REV7 antibody (612266, BD Biosciences, 1:200 dilution) and anti- γ H2AX or anti-53BP1 in 5% goat serum for 1 h at room temperature. Coverslips were washed and incubated with secondary antibodies conjugated with rhodamine or FITC for 30 min at room temperature. Cells were then stained with DAPI to visualize nuclear DNA. The coverslips were mounted onto glass slides with anti-fade solution. For RIF1 staining in MEF cells, wild-type MEFs stably transduced with indicated shRNA-expressing lentiviruses were examined for RIF1 foci following neocarzinostatin treatment. Automated quantification of RIF1 foci following mock and neocarzinostatin treatment was performed using Cell-Profiler software (Broad Institute).

Laser irradiation of human cells and immunofluorescence staining. Cells were grown on plastic disks (17 mm diameter) that were cut using CNC cutter from the bottom of standard 10-cm cultivation dish (TPP) ultraviolet-sterilized and placed inside the wells of a 12-well plate. BrdU (0.5 μ M) was added into siRNA-transfected cells 40 h after the transfection to pre-sensitize cells towards UV-A wavelength. Twenty-four hours after BrdU addition, the plastic disks with cells were removed and covered by a coverslip and immediately placed inside Zeiss AxioObserver Z.1 inverted microscope combined with LSM 780 confocal module. Cells were irradiated at 20 °C via \times 40 water immersion objective (Zeiss C-Apo 40 \times /1.2W DICIII), using 355 nm 65 mW laser set on 100% power. The total laser dose that can be further manipulated by the amount of irradiation cycles was empirically set to six irradiation cycles. Laser track was pre-defined to cover all the cells within the acquisition area with at least one stripe across the nucleus. After the irradiation process the coverslip was gently removed and plastic disk was quickly placed back into the same well of the 12-well plate and incubated for another 45 min at standard cultivation conditions. The plastic disks with laser-irradiated cells were first processed by pre-extraction at 4 °C. It involves washing by PBS (4 °C), equilibration for 2 min in sucrose buffer 1 (10 mM PIPES, pH 6.8, 100 mM NaCl, 1.5 mM MgCl₂ and 300 mM sucrose) on ice and then pre-extraction for 15 min on ice, on slow-moving shaker using sucrose buffer 2 (10 mM PIPES, pH 6.8, 100 mM NaCl, 1.5 mM MgCl₂, 300 mM sucrose, 0.5% Triton X-100, 5 μ g ml⁻¹ leupeptin, 2 μ g ml⁻¹ aprotinin, 0.1 mM phenylmethanesulfonylfluoride (PMSF), 1 mM dithiothreitol (DTT)). After the pre-extraction cells were washed by PBS (4 °C) and fixed by 4% paraformaldehyde (PFA) for 15 min at room temperature. PFA was washed out three times by PBS. The disks were further processed as standard coverslips (that is, blocking in blocking solution (DMEM plus 20% FCS) for 1 h followed by incubation with primary antibodies involving REV7 (BD Bioscience, mouse, 1:200), pS139-H2AX (Cell Signaling, 20E3, rabbit, 1:300) and 53BP1 (Santa Cruz, H-300, rabbit 1:400) for 2 h, and with appropriate secondary antibodies coupled to Alexa Fluor 488 and Alexa Fluor 568 fluorophores (dilution 1:1,000) (Life Technologies). Both primary and secondary antibodies were dissolved in the blocking solution. After washes in PBS, the disks were incubated in 1 μ g ml⁻¹ DAPI in dH₂O at room temperature for 5 min and air dried. Dried disks were placed on a standard microscopy glass (cell layer face up) and anchored by two rubber bands laced over the glass. Stained cells were mounted using VectaShield (Vector Labs) mounting medium and covered by a coverslip. The samples were examined using Zeiss Axio-Observer Z.1 inverted microscope combined with LSM 780 confocal module using

\times 40 oil objective (Zeiss EC PlN 40 \times /1.3 Oil DICII). It means that after the first acquisition the plastic disk and the microscopy glass was marked by diamond cutter (to ensure same positioning of the disk in the future), the coverslip was gently removed and disk was washed three times in PBS, 0.5% Tween to remove the mounting medium. Next, the disk was incubated in the 1 \times Re-Blot solution (Re-Blot Plus Mild, Millipore) for 30 min on a slow moving shaker. The solution was washed out three times in PBS. Such sample was ready for new staining procedure involving new set of primary and secondary antibodies following the same protocol as described above.

In situ analysis of RAD51 foci formation. Five matched PARPi-resistant and -sensitive KB1P(M) tumours were orthotopically transplanted into wild-type FVB/N recipient mice. When tumours reached \sim 500 mm³ in volume, the mice were randomized to be either irradiated (dose: 15 Gy) using a CT-guided high precision cone beam micro-irradiator (X-RAD 225Cx) or left untreated. As a positive control a BRCA1-proficient KP tumour was taken along. Two hours after irradiation the tumours were taken out and fixed in 4% formalin. Immunofluorescence staining was performed on FFPE slides. RAD51 foci were detected using a non-commercial antibody provided by R. Kanaar in a dilution of 1:5,000. 53BP1 foci were detected using rabbit anti-53BP1 (A330-272A, Bethyl), diluted 1:500. As a secondary antibody goat anti-rabbit-Alexa Fluor 568 (Invitrogen) was used at a dilution of 1:1,000 (2 μ g ml⁻¹). Images were taken by a 'blinded' investigator using a confocal microscope (Leica SP5, Leica Microsystems GmbH), equipped with a \times 100 objective. For each tumour five random areas (246 \times 246 μ m) were imaged. Image stacks (\sim four slices) were analysed in ImageJ, using an in-house developed macro to automatically and objectively evaluate the RAD51 foci. In brief, nuclei were segmented by thresholding the (median-filtered) DAPI signal, followed by a watershed operation to separate touching nuclei. For every z-stack the maximum-intensity projection of the foci signal was background-subtracted using a difference of gaussians method. Next, for each nucleus, foci candidates were identified as locations where the resulting pixel values exceeded the background by a factor (typically tenfold) times the median standard deviation of all nuclei in the image. In combination with additional filters discriminating for foci size and absolute brightness this procedure yielded a robust and reliable foci count for all nuclei. Results were validated by visual inspection.

REV7 recruitment to local laser-induced DNA damage sites. pEGFP-REV7 or pEGFP and 53BP1-mCherry were co-transfected into *Brcal*^{-/-} *p53*^{-/-} cells using X-treme GENE HP DNA Transfection Reagent according to the manufacturer's manual. GFP and mCherry double-positive cells were sorted by flow cytometry and seeded onto coverslips. Cells were sensitized by pre-incubation with Hoechst33342 and were subsequently irradiated using a 405-nm diode laser (\times 63 objective, 0.99 mW, 60% laser power, 50 s) on a Leica SP5 confocal microscope equipped for live-cell imaging. EGFP-REV7 and 53BP1-mCherry recruitment in living cells was monitored by time-lapse imaging.

Alpha track assay. Cells were seeded in dishes with a mylar surface as previously described²⁸, allowing α -particle irradiation through the bottom of the dish. One or two hours after irradiation three times for 30 s with a ²⁴¹americium source, cells were washed once in ice-cold PBS. Subsequently, cells were extracted with cold CSK buffer (10 mM HEPES-KOH, pH 7.9, 100 mM NaCl, 300 mM sucrose, 3 mM MgCl₂, 1 mM EGTA, 0.5% (v/v) Triton X-100) and cold CSS buffer (10 mM Tris, pH 7.4, 10 mM NaCl, 3 mM MgCl₂, 1% (v/v) Tween20, 0.5% (w/v) sodium deoxycholate) for 5 min each before fixation in 4% PFA in PBS for 20 min at room temperature. Fixed cells were washed five times in PBS plus 0.1% Triton X-100 and washed once in blocking solution (0.5% BSA plus 0.15% glycine in PBS). Primary antibodies were diluted in blocking solution and cells were incubated overnight at 4 °C. After incubation, cells were washed five times with PBS plus 0.1% Triton X-100 and washed once in blocking solution. Secondary antibodies were diluted in blocking solution and cells were incubated at room temperature for at least 1 h. Afterwards, cells were washed five times in PBS plus 0.1% Triton X-100 and once in PBS. Finally, mylar films were glued on glass slides and cells were mounted using Vectashield with DAPI. For quantification, at least 100 53BP1 or MRE11-positive tracks were scored for the presence of RPA. Primary antibodies used in this study were as follows: rabbit anti-53BP1 (NB100-304, Novus), 1:1,000 dilution; mouse anti-RPA2 (Ab2175, Abcam), 1:500 dilution; MRE11 antibody²⁹, 1:200 dilution. Secondary antibodies used in this study were as follows: Alexa Fluor 594 goat anti-rabbit IgG (A 31631, Invitrogen), Alexa Fluor 488 goat anti-mouse IgG (A11001, Invitrogen).

BrdU propidium iodide cell cycle assay. Cells were seeded in 6-cm dishes and attached overnight. The next day, cells were incubated for 15 min with 5 μ M BrdU in growth medium, trypsinized, washed in PBS and fixed in 70% ethanol overnight. Fixed cells were washed in PBS, resuspended in pepsin solution (5 mg pepsin in 10 ml 0.1 N HCl) and incubated for 20 min at room temperature. Subsequently blocking solution (0.5% Tween, 0.1% BSA in PBS) was added to wash. Next, cells were resuspended in 2 N HCl and incubated 12 min at 37 °C. To neutralize, 100 mM

borate buffer (pH 8.5) was added and cells were pelleted. Anti-BrdU-FITC antibody (347583, BD Bioscience) was diluted in blocking solution, added to cells and incubated for 2 h on ice. Stained cells were washed once with blocking solution and resuspended in 500 µl PBS plus 12.5 µl RNaseA and 1 µl propidium iodide (P3566, Invitrogen). Cell cycle analysis was performed the next day.

Survival assay of mES cells. R26^{CreERT2/wt}; *Brcal*^{SCo/A} mES cells¹⁵ were infected with hairpins targeting *Rev7* or the vector control and selected with puromycin. Expression of mouse *Brcal* was switched off by overnight incubation with 0.5 µM 4-hydroxytamoxifen. Four days after switching, 5,000 cells of the indicated groups were seeded per well into 6-well plates and assayed for growth. Surviving colonies were fixed in formalin and stained with crystal violet.

DR-GFP for HR assay. R26^{CreERT2/wt}; *Brcal*^{SCo/A} mES cells were targeted with a modified version of the p59X DR-GFP construct as described³⁰. To allow experiments on a p53-deficient background, cells were infected with a lentiviral p53 shRNA (5'-AGAGTATTCACCCCTCAAGAT-3') using a pLKO1 vector provided with a neomycin resistance marker. A G418-selected p53-deficient clone was subsequently infected with hairpins targeting *Rev7* or the vector control and selected with puromycin. Expression of mouse *Brcal* was switched off by overnight incubation with 0.5 µM 4-hydroxytamoxifen to measure the effect of *Rev7* loss on HR. HR reporter assays were performed by Lipofectamine 2000 (Invitrogen) transfection of the I-SceI-mCherry plasmid, which was generated by providing the cBasI-SceI expression plasmid with CMV-mCherry (Clontech). Forty-eight hours after transfection, mCherry/GFP double-positive cells were monitored by flow cytometry as described¹⁵.

53BP1 pull-down. 53BP1 pull-downs were performed as described¹⁹. Flag pull-downs were performed from 2 mg lysate prepared from MEFs following mock or neocarzinostatin treatment (2 h at 250 ng µl⁻¹). Control, 53BP1 and 53BP1^{20AQ} immunoprecipitates were then treated with sequential low-salt (150 mM) and high-salt (500 mM) RIPA buffer washes, before re-equilibration in nuclear extract buffer (20 mM HEPES, pH 7.9, 100 mM KCl, 0.2 mM EDTA, 20% glycerol, 0.5 mM PMSF, 0.5 mM DTT and protease inhibitors (Roche, Complete)). After incubation in HeLa Nuclear Extract (2 mg), control, 53BP1 and 53BP1^{20AQ} complexes were washed four times in nuclear extract buffer, then eluted with triple-Flag peptide (Sigma).

Immunoglobulin CSR.

CH12 cells were either mock-treated or stimulated with agonist anti-CD40 antibody (0.5 µg ml⁻¹; eBioscience; HM40-3), mouse IL-4 (5 ng µl⁻¹; R&D Systems) and TGF-β1 (1.25 ng µl⁻¹; R&D Systems). Cell-surface IgA expression was determined by flow cytometry, immunostaining with biotinylated anti-mouse IgA antibody (eBioscience; 13-5994), and Alexa488-streptavidin conjugate (Life Technologies).

Carboxyfluorescein succinimidyl ester assay. Cell proliferation was assessed in stimulated CH12 cells using carboxyfluorescein succinimidyl ester (CFSE) according to manufacturer's instructions (CellTrace; Life Technologies).

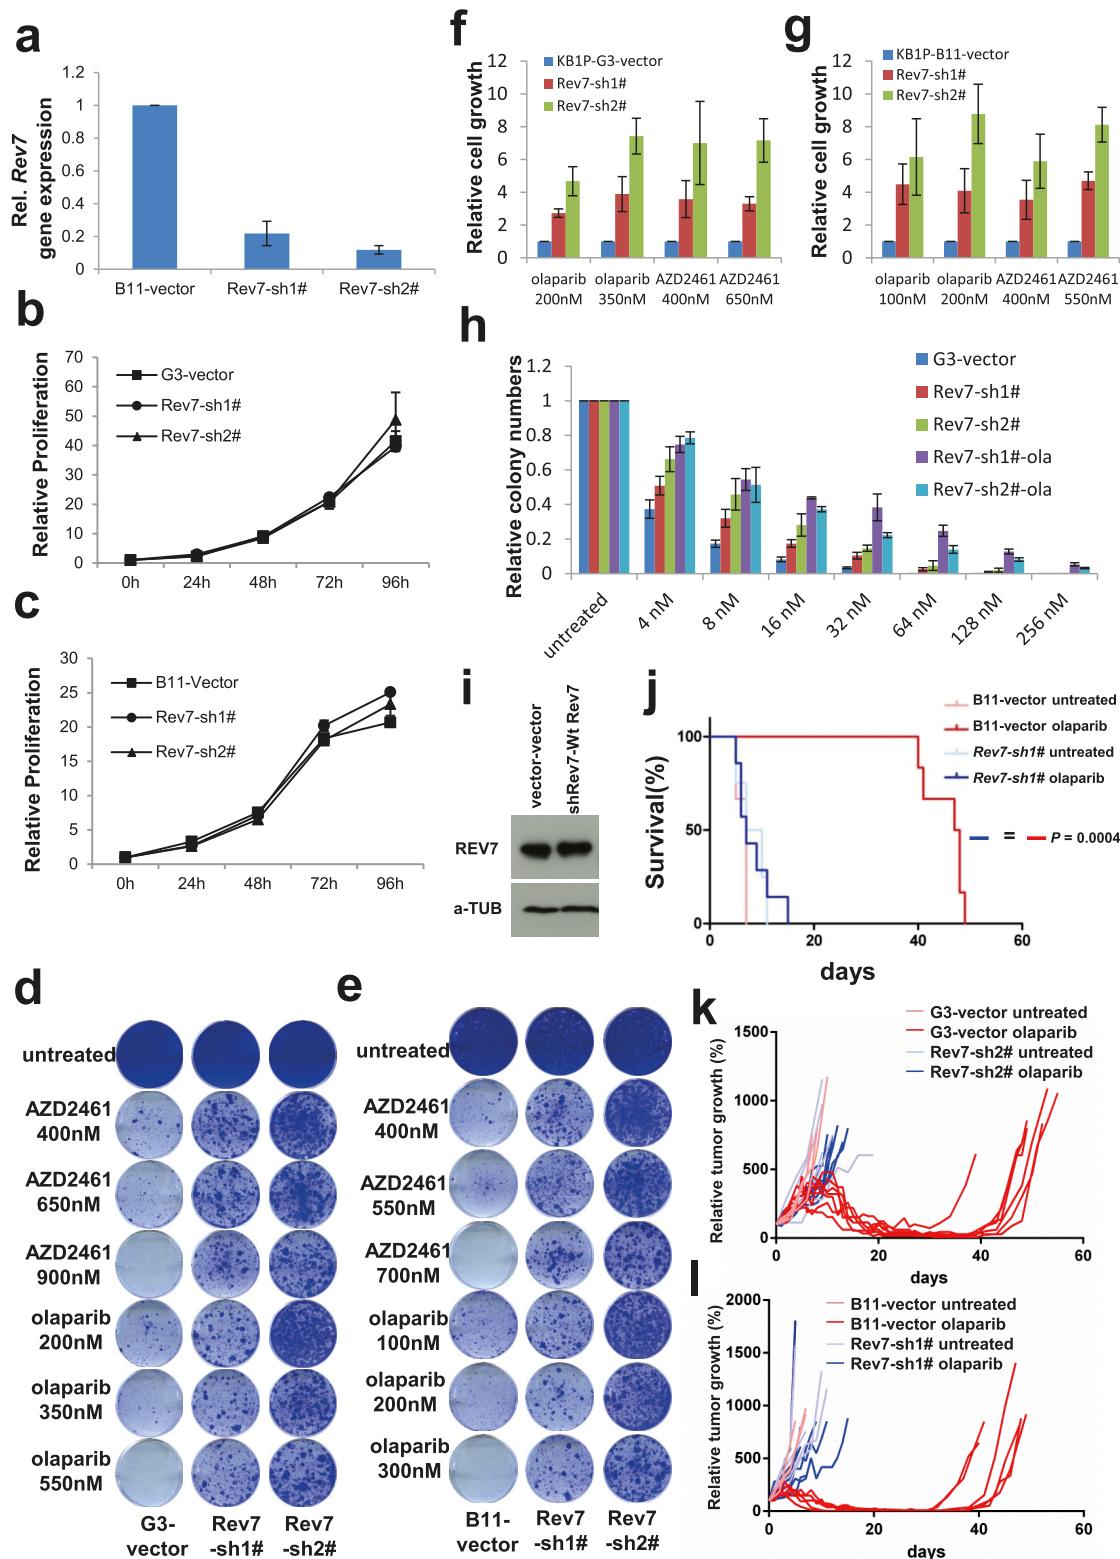
Chromatin immunoprecipitation. Each ChIP was performed from chromatin prepared from ~1 × 10⁷ CH12 cells stimulated for 30 h with agonist CD40 antibody, IL-4 and TGF-β1 essentially as described¹⁹. For each individual ChIP, 4 µg of RPA34-20 (Ab-3, Calbiochem), 2 µg H2AX (3522-1, Epitomics), or 4 µg control mouse anti-IgG (sc-2025; Santa Cruz) coupled to 25 µl Protein-G Dynabeads (Life Technologies, 10003D) was used. Quantities of immunoprecipitated chromatin were calculated relative to total input chromatin by quantitative PCR in duplicate on an CFX96 Real-Time Analyzer (Biorad) with the use of iQ SYBR Green (Biorad) for each primer pair (see below). *Rpp30*: forward, 5'-TCCAGTGTGCAAG AAAGCTAAATG-3'; reverse, 5'-GGCAGTGCCTGGAGACTCA-3'; A (target IgH Sp1): forward, 5'-CAATGTGGTTAACGTTGAAGTTGCCA-3'; reverse, 5'-TCTCACACTCACCTGGATCTAACGACTGT-3'; B (target IgH Sp1): forward, 5'-GCTAAACTGAGGTGATTACTCTGAGGTAAAG-3'; reverse, 5'-GTT TAGCTTAGCGGCCAGCTCATTCCAGT-3'; C (target IgH Sp1): forward, 5'-AGTGTGGAAACCCAGTCAAA-3'; reverse, 5'-GTACTCTCACCGGGATCA GC-3'; D (target IgH Sp1): forward, 5'-TGAAAAGACTTGGATGAAATGTGA ACCAA-3'; reverse, 5'-GATACTAGGTGATGGCTCCATTACACACA-3'.

Immunohistochemistry on paraffin sections. The panel of formalin-fixed, paraffin-embedded archival specimens from a series of 50 human primary breast carcinomas

was examined, surgically resected before any chemotherapy or radiotherapy. In parallel, sections from ten specimens of formalin-fixed, paraffin-embedded histologically normal human breast tissues were used as normal tissue controls, for immunohistochemical analysis of *REV7* expression patterns. The breast carcinomas were all of the triple-negative type, defined in this cohort as having fewer than 1% tumour cells positive in standard immunohistochemical staining for oestrogen and progesterone receptor proteins, and lacking amplification of the *HER2* (also known as *ERBB2*) gene. For the immunohistochemical analysis, 4-µm-thick sections were cut from representative blocks of the tumour tissues, the sections were deparaffinized and processed for sensitive immunoperoxidase staining with the primary mouse monoclonal antibody to *REV7* (BD Biosciences, 612266, diluted 1:100). The staining procedure was essentially as described¹⁵, with antigen unmasking in citrate buffer, pH 6.0, for 15–20 min in a microwave oven, and overnight incubation with the primary antibody, followed by the Vectastain Elite kit (Vector Laboratories) and nickel sulfate enhancement without nuclear counterstaining³¹. Mouse normal serum and antibody to γH2AX served as negative and positive controls, respectively. Evaluation of the staining patterns was performed by two independent observers (with very similar outcomes), including a senior oncopathologist with more than 20 years of experience with breast cancer pathology. The results were scored in the following categories, based on comparison of cancer cells with the series of normal breast tissue controls, and also the normal of cells present within each of the tumour sections. As normal breast epithelium showed reproducible positivity for *REV7* protein in over 90% of cells, we considered as aberrantly decreased expression in a tumour when fewer than 70% of cancer cells were positive. In addition to the percentage of stained tumour cells, staining intensity was classified as either normal (comparable with the intensity of normal cells present on each section) or aberrantly low (clearly below the intensity seen in adjacent normal cells, and up to undetectable in some cases). Overall, while 6 out of 47 informative cases showed concomitantly aberrant fraction of *REV7*-stained cells and reduced intensity of staining, 12 additional cases showed less pronounced defects limited to either staining intensity or reduced percentage of cancer cells, respectively. As the primary goal of these analyses was to establish a detection assay for *REV7* on archival tissue specimens and to assess the frequency of potentially *REV7*-deficient breast tumours, correlation analyses with clinical parameters including treatment outcome remain to be performed on larger cohorts of patients in the future.

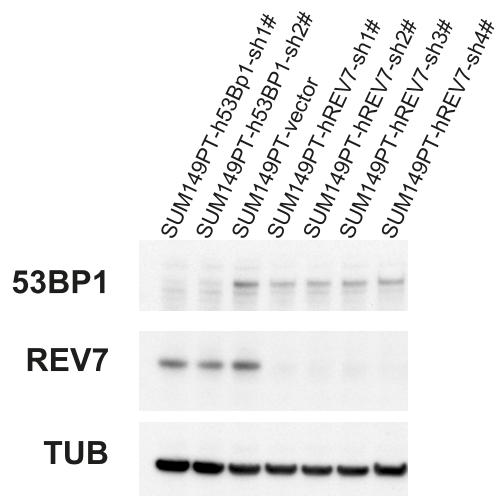
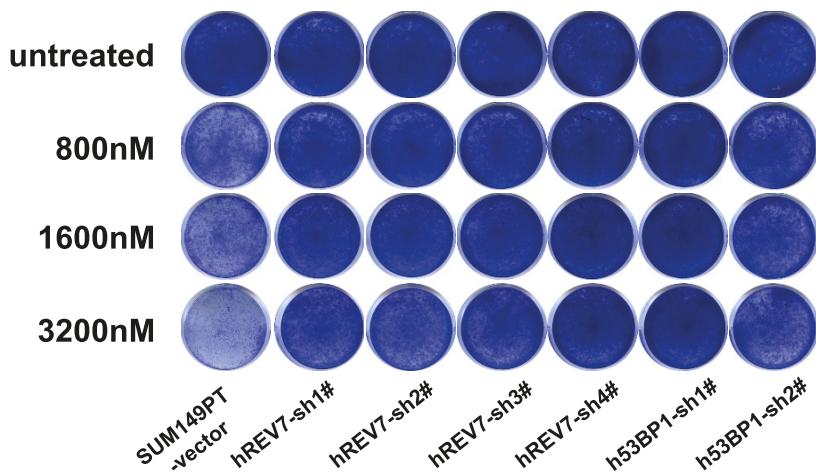
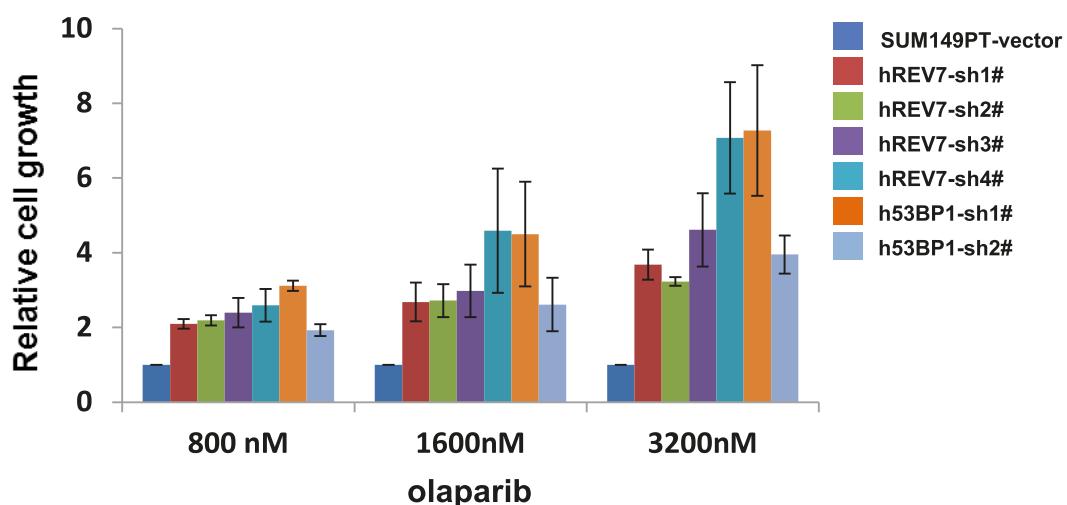
Statistics. Statistical tests used were log-rank test, *t*-test and Mann–Whitney *U* test, with *P* < 0.05 as the significance level. No statistical methods were used to pre-determine sample size.

23. Evers, B. et al. Selective inhibition of BRCA2-deficient mammary tumor cell growth by AZD2281 and cisplatin. *Clin. Cancer Res.* **14**, 3916–3925 (2008).
24. Fradet-Turcotte, A. et al. 53BP1 is a reader of the DNA-damage-induced H2A Lys 15 ubiquitin mark. *Nature* **499**, 50–54 (2013).
25. Watanabe, S. et al. JMJD1C demethylates MDC1 to regulate the RNF8 and BRCA1-mediated chromatin response to DNA breaks. *Nature Struct. Mol. Biol.* **20**, 1425–1433 (2013).
26. Doil, C. et al. RNF168 binds and amplifies ubiquitin conjugates on damaged chromosomes to allow accumulation of repair proteins. *Cell* **136**, 435–446 (2009).
27. Stewart, G. S. et al. The RIDDLE syndrome protein mediates a ubiquitin-dependent signaling cascade at sites of DNA damage. *Cell* **136**, 420–434 (2009).
28. Stap, J. et al. Induction of linear tracks of DNA double-strand breaks by α-particle irradiation of cells. *Nature Methods* **5**, 261–266 (2008).
29. de Jager, M. et al. DNA-binding and strand-annealing activities of human Mre11: implications for its roles in DNA double-strand break repair pathways. *Nucleic Acids Res.* **29**, 1317–1325 (2001).
30. Bouwman, P. et al. A high-throughput functional complementation assay for classification of BRCA1 missense variants. *Cancer Discov.* **3**, 1142–1155 (2013).
31. Bartkova, J. et al. DNA damage response as a candidate anti-cancer barrier in early human tumorigenesis. *Nature* **434**, 864–870 (2005).



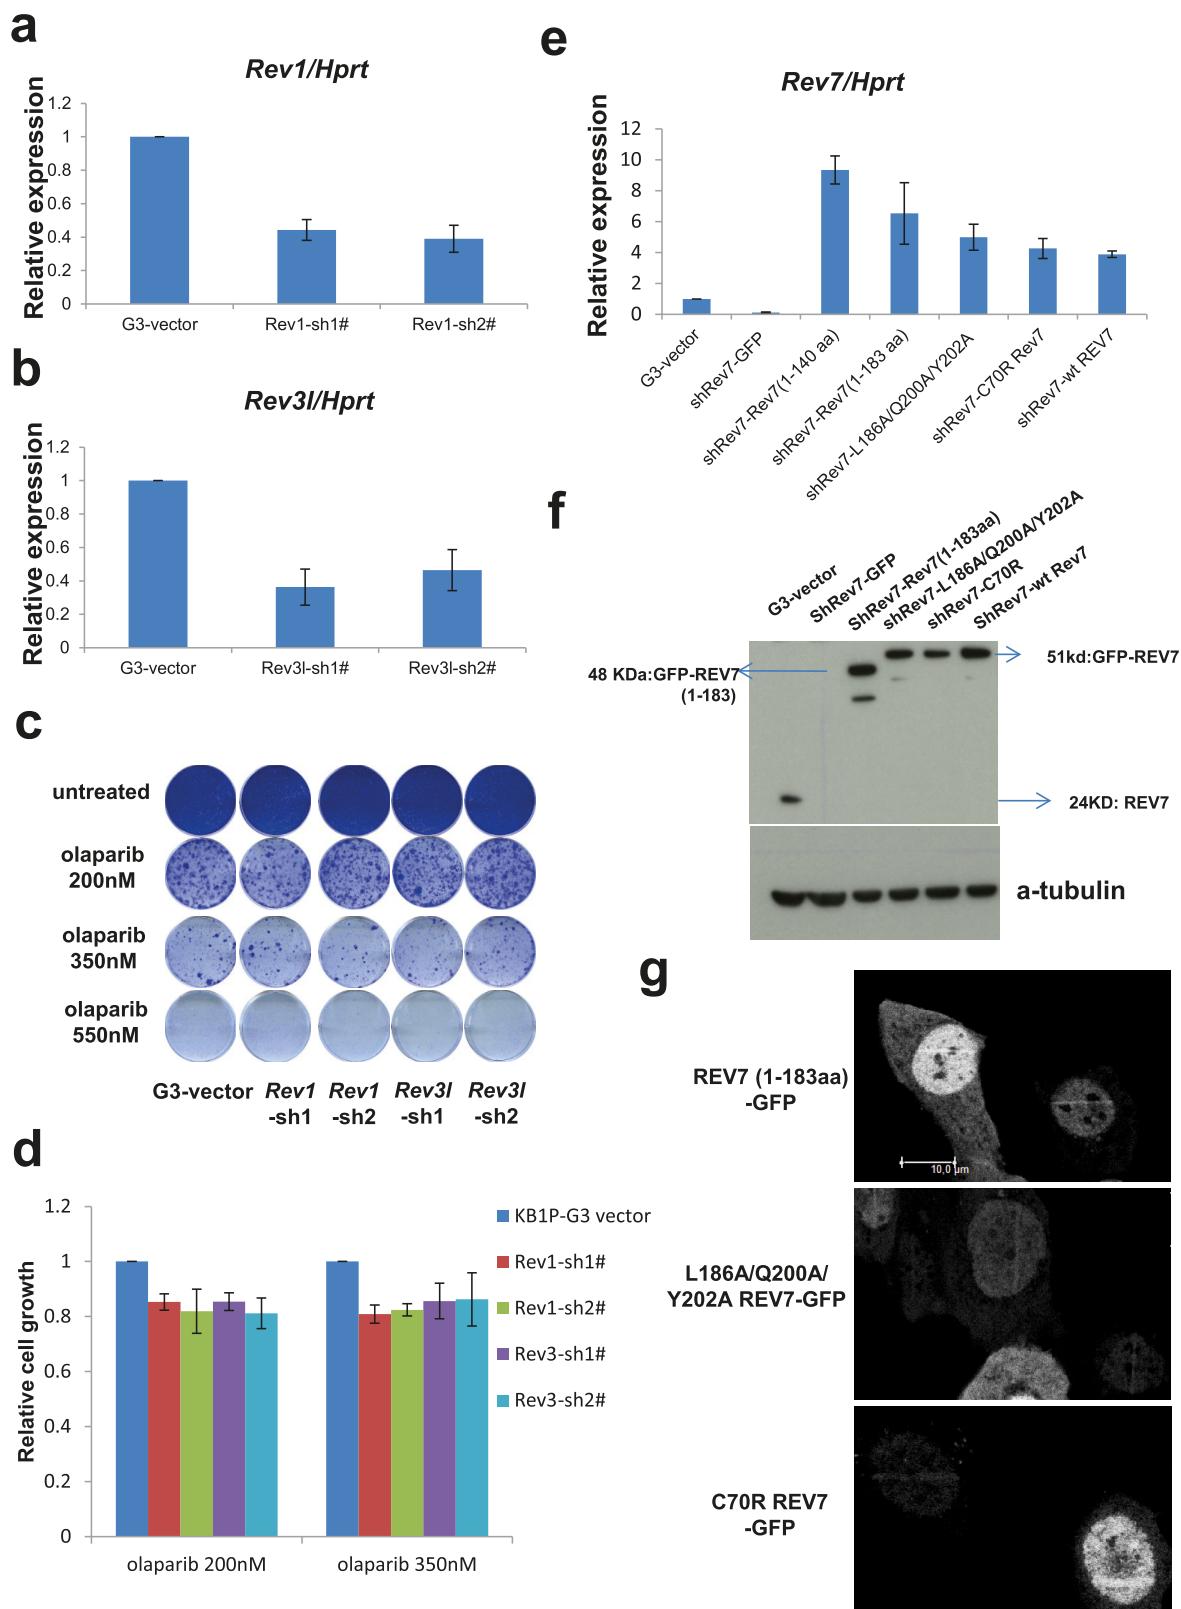
Extended Data Figure 1 | Loss of Rev7 causes PARPi resistance in vitro and in vivo. **a**, Quantification of Rev7 transcript levels in KB1P-B11 cells transduced with Rev7-targeting shRNAs or the vector control. *Hprt* was used as a control for transcript expression. The data represent mean \pm s.d. **b, c**, Cell proliferation rates in KB1P-G3 (**b**) or KB1P-B11 (**c**) cells analysed using the MTT assay. **d–g**, Long-term clonogenic survival assays and quantification of KB1P-G3 (**d, f**) or KB1P-B11 (**e, g**) cells transduced with the indicated constructs and treatments. All the groups were normalized by the absorbance of the vector control. The data represent mean \pm s.d. **h**, Quantification of the real colony numbers from the short-term clonogenic survival assay of KB1P-G3

cells with or without Rev7 loss exposed to olaparib. **i**, REV7 protein levels were determined by western blotting of lysates derived from KB1P-G3 cells transduced with the indicated constructs. **j**, Overall survival of mice with KB1P-B11-derived Rev7-depleted or control tumours treated with one regimen of 50 mg kg^{-1} olaparib daily for 28 days or left untreated. The *P* value was calculated using the log-rank test. **k, l**, Relative tumour growth of individual KB1P-G3- (**k**) and KB1P-B11- (**l**) derived Rev7-depleted or control tumours treated with one regimen of 50 mg kg^{-1} olaparib daily for 28 days or left untreated.

a**b****c**

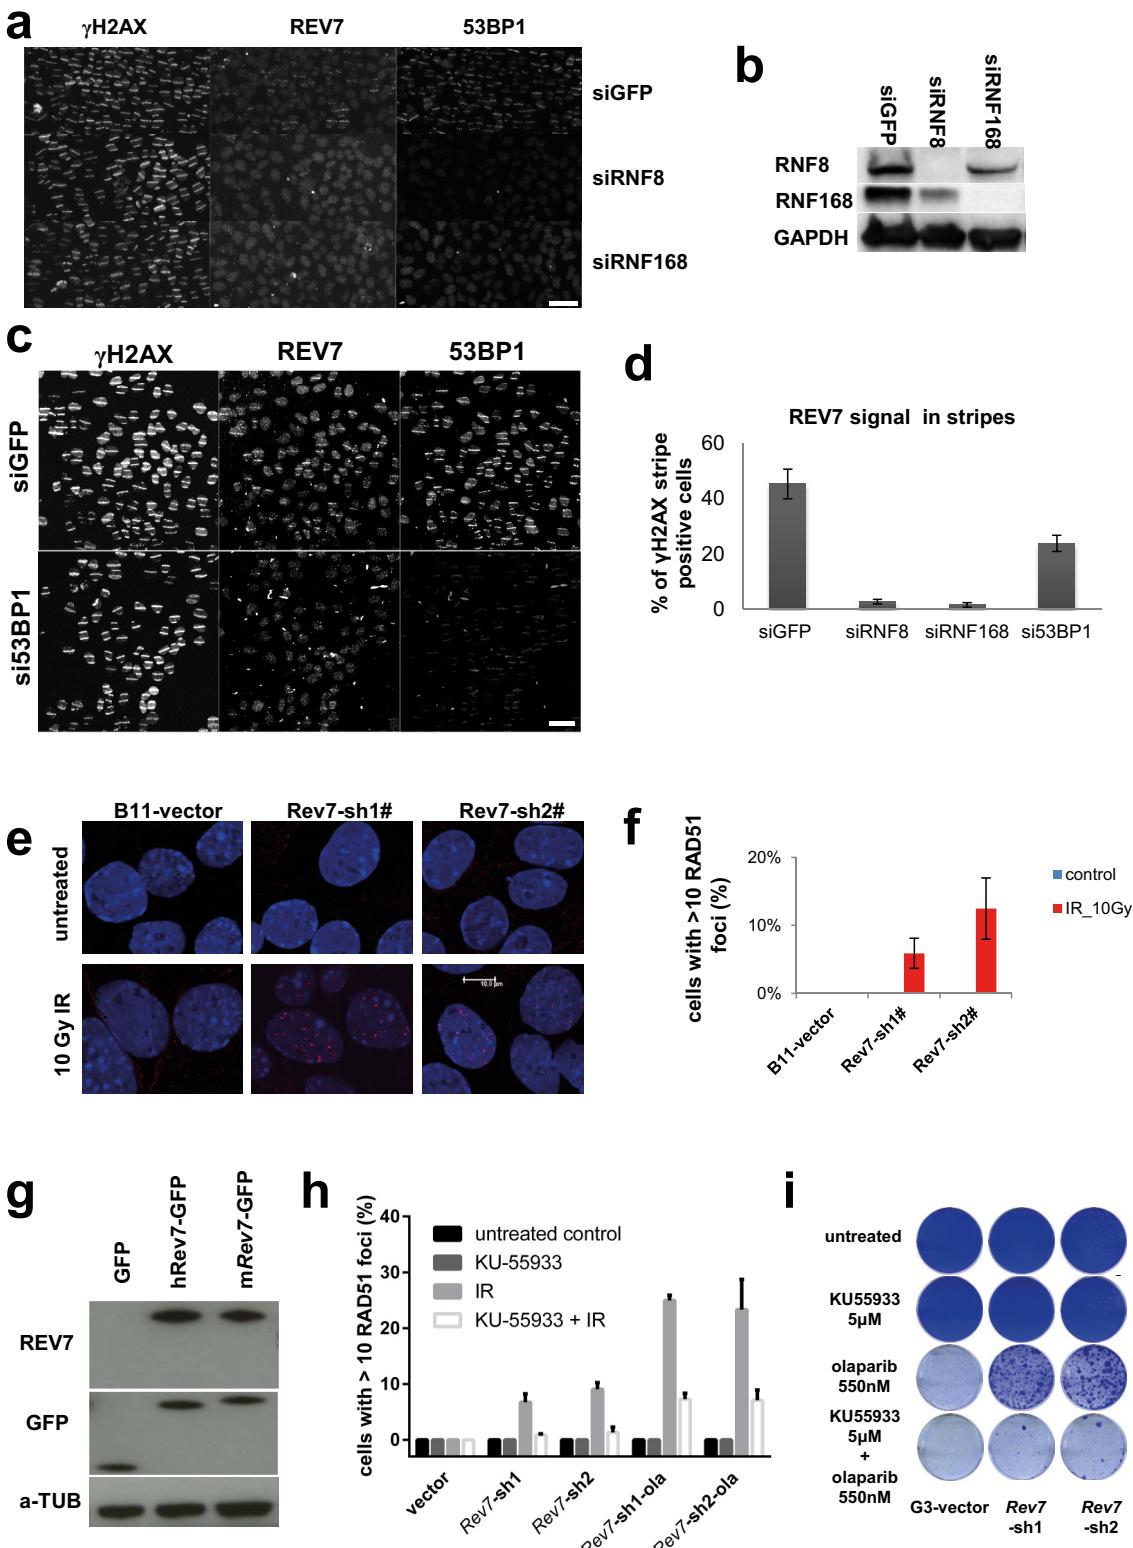
Extended Data Figure 2 | Loss of REV7 causes olaparib resistance in BRCA1-deficient SUM149PT cells. **a**, Western blotting analysis of REV7 or 53BP1 expression in SUM149PT cells transduced with REV7- or 53BP1-targeting hairpins or the vector control. **b**, Example of a long-term clonogenic survival assay using the indicated hairpins and olaparib concentrations.

c, Quantification of the clonogenic assays using absorbance of crystal violet at 590 nm. The data represent mean \pm s.d. All the groups were normalized by the absorbance of the vector control and showed significant differences to the control ($P < 0.01$, *t*-test).



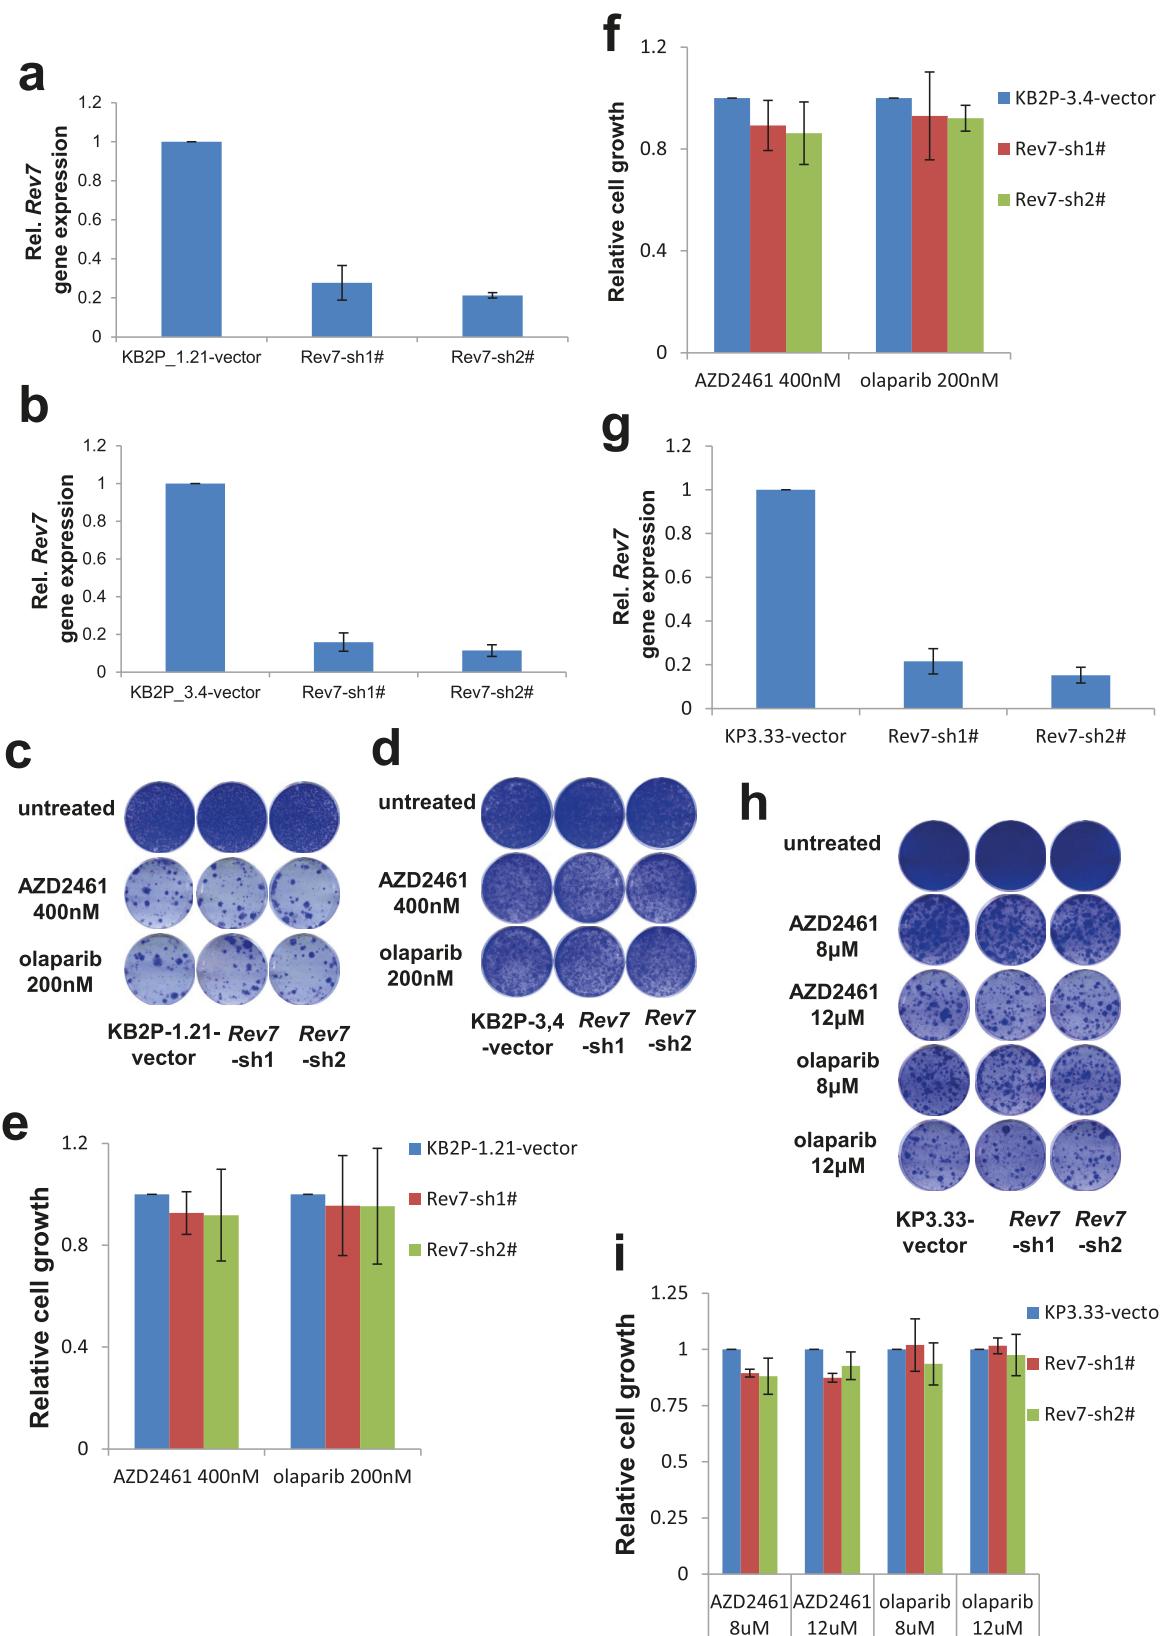
Extended Data Figure 3 | *Rev1* or *Rev3* inhibition and PARPi sensitivity of *Brcal*^{-/-} *p53*^{-/-} mammary tumour cells. **a, b,** Quantification of *Rev1* (**a**) or *Rev3* (**b**) transcript levels in KB1P-G3 cells transduced with *Rev1*- or *Rev3*-targeting shRNAs or the vector control. *Hprt* was used as a control for transcript expression. The data represent the mean \pm s.d. **c,** Long-term clonogenic survival assays of KB1P-G3 cells exposed to the indicated PARP inhibitors. **d,** Quantification of the clonogenic assay by determining the absorbance of crystal violet at 590 nm. All the groups were normalized by the

absorbance of the vector control. The data represent the mean \pm s.d. **e, f,** Quantification of *Rev7* transcript (**e**) or protein (**f**) levels in KB1P-G3 cells transduced with *Rev7*-targeting shRNAs or the vector control. *Hprt* was used as a control for transcript expression; α -tubulin as a control for protein expression. The data represent mean \pm s.d. **g,** GFP-tagged *REV7* mutants recruitment to sites of DNA damage was observed 5 min after 405 nm laser exposure (0.99 mW, 60% laser power, 50 s) in KB1P-B11 cells. Scale bar, 10 μ m.



Extended Data Figure 4 | REV7 recruitment to the DNA damage sites in human cells. a–c, Human REV7 recruitment to sites of laser-induced DNA damage was analysed in U2OS cells transfected with siRNAs targeting RNF8 (siRNF8), RNF168 (siRNF168) (a), 53BP1 (si53BP1) (c) and GFP (siGFP). RNF8 and RNF168 protein levels were determined by western blotting (b) using lysates derived from U2OS cells transfecting with the indicated siRNAs. d, For the quantification of the REV7 signal within laser-induced DNA damage stripes, a minimum of 100 striped (that is, γ H2AX-positive) cells were analysed for the presence of the REV7 signal in two independent experiments. Scale bars, 50 μ m. e, RAD51 focus (red) formation in KB1P-B11 cells.

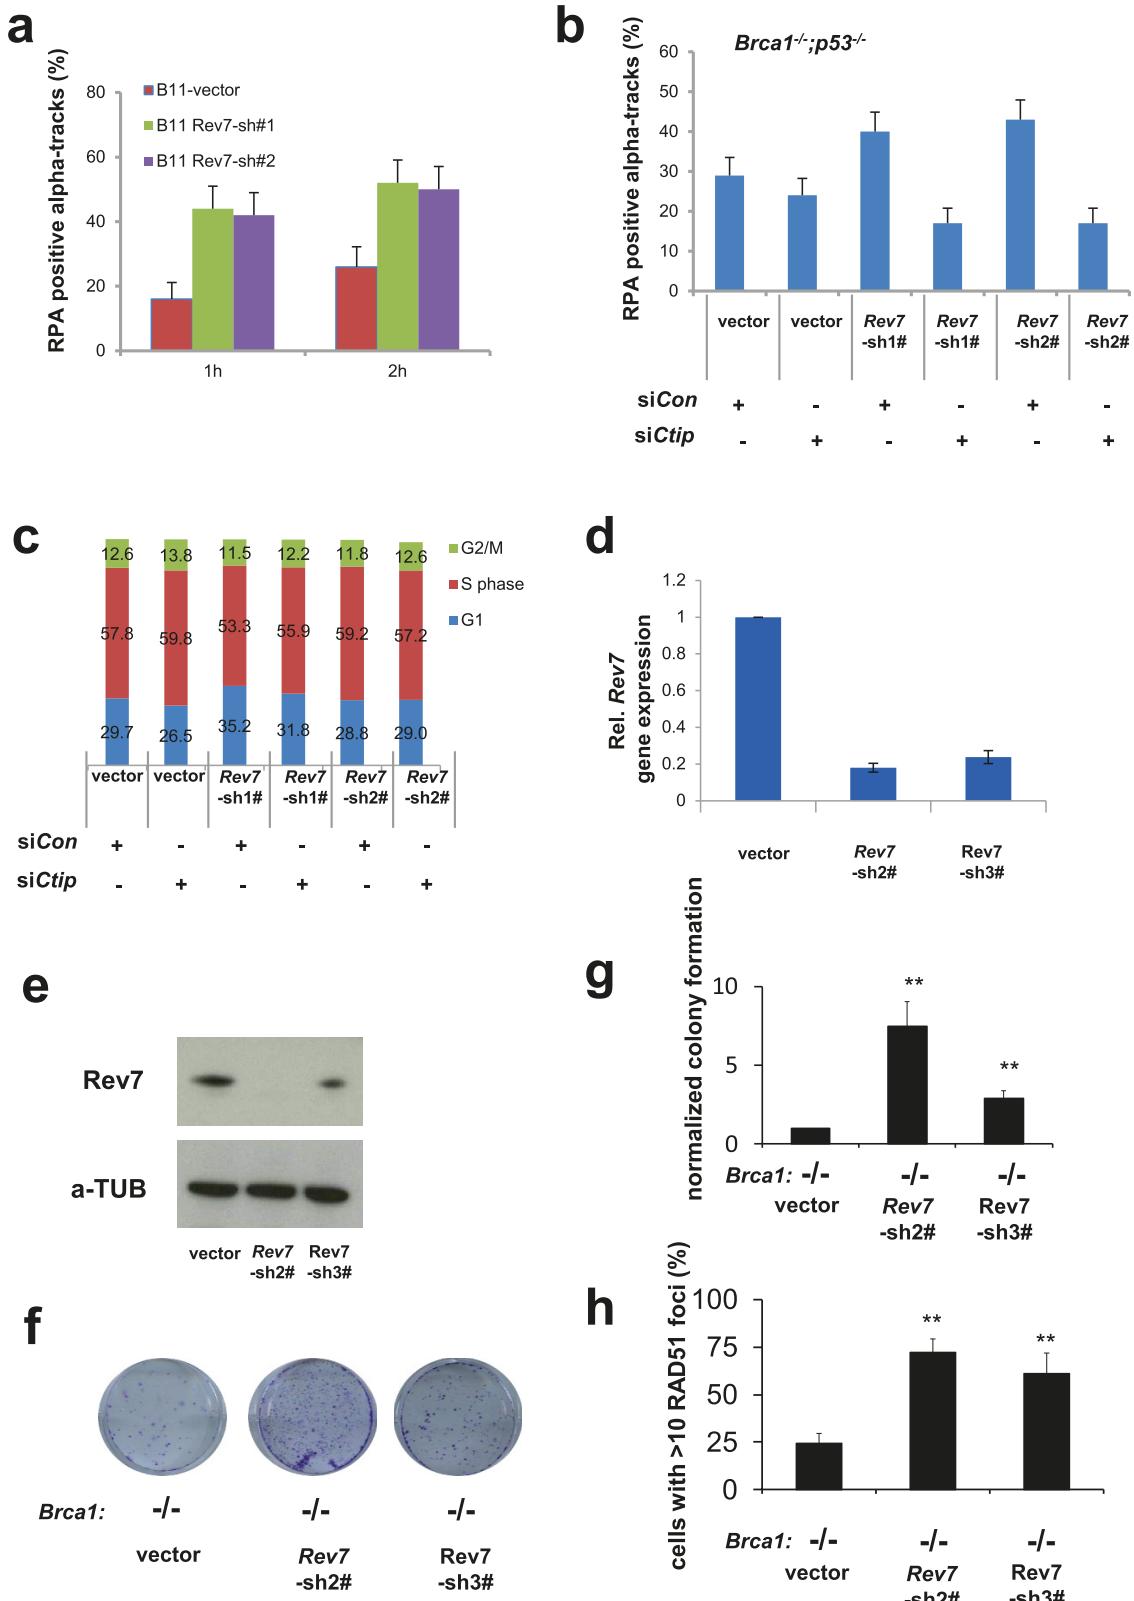
cells before and 5 h after 10 Gy ionizing radiation. Scale bar, 10 μ m. f, Quantification of RAD51 foci in KB1P-B11 cells in the presence or absence of REV7 depletion. At least 150 cells were analysed per group in three independent experiments each. Error bars indicate s.d.; IR denotes 5 h after 10 Gy ionizing radiation. g, Western blotting analysis of REV7-depleted KB1P-G3 cells transfected with human REV7-GFP or Rev7-shRNA-resistant mouse REV7-GFP fusion proteins. h, Same as in e and f using the ATM inhibitor KU55933 with or without IR (5 h after 10 Gy ionizing radiation). i, Long-term clonogenic survival assay of KB1P-G3 cells exposed to olaparib in the presence or absence of KU55933 pre-treatment.



Extended Data Figure 5 | Loss of Rev7 does not cause PARPi resistance in *Brca2*^{-/-} *p53*^{-/-} or *p53*^{-/-} mammary tumour cells *in vitro*.

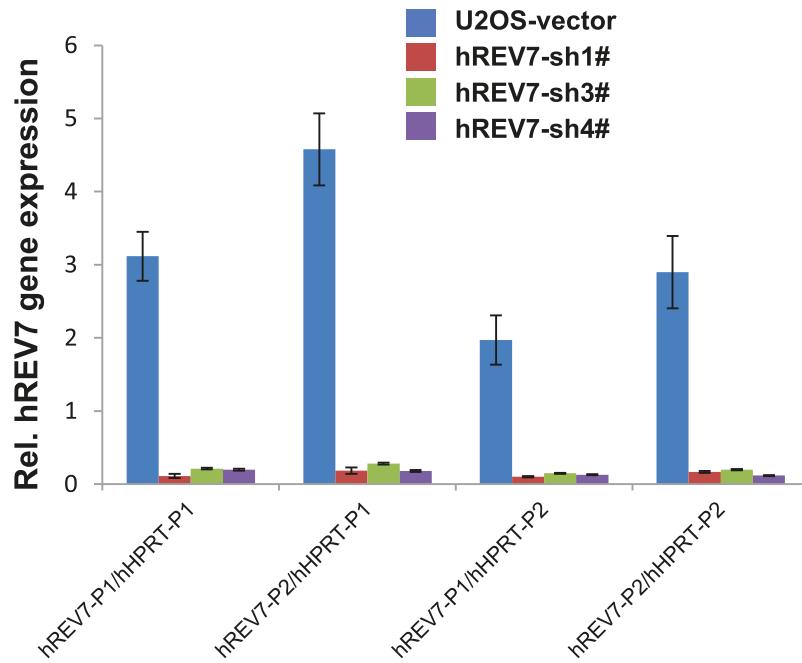
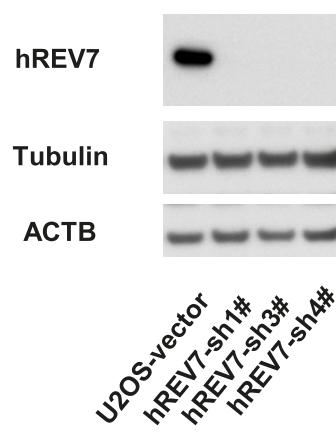
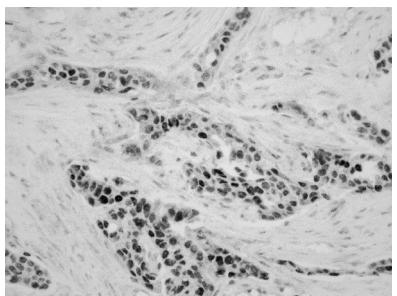
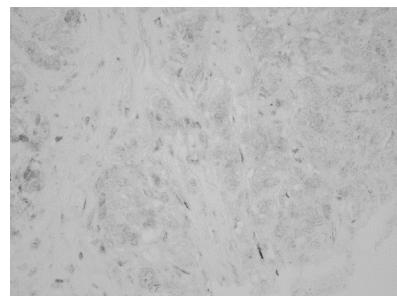
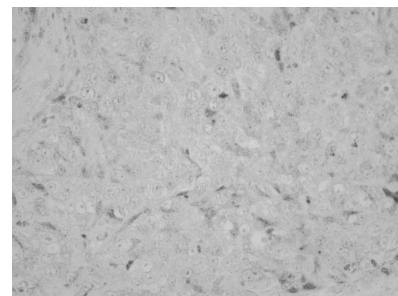
a, b, Quantification of Rev7 transcript levels in *Brca2*^{-/-} *p53*^{-/-} (KB2P-1.21 or KB2P-3.4) cells transduced with Rev7-targeting shRNAs or the vector control. *Hprt* was used as a control for transcript expression. The data represent the mean \pm s.d. c–f, Long-term clonogenic survival assays and quantification of KB2P-1.21 or KB2P-3.4 cells with or without Rev7 depletion exposed to the

indicated treatments. All the groups were normalized by the absorbance of the vector control. The data represent mean \pm s.d. g, Quantification of Rev7 transcript levels in *p53*^{-/-} (KP3.33) cells transduced with the indicated constructs. *Hprt* was used as a control for transcript expression and the data represent the mean \pm s.d. h, i, Long-term clonogenic survival assays and quantification of KP3.33 cells exposed to the indicated treatments. The data represent the mean \pm s.d.



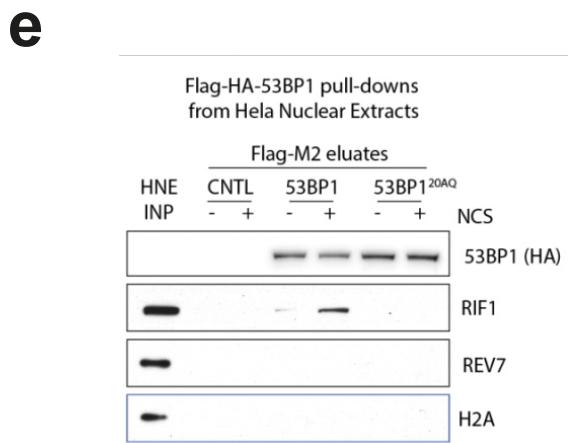
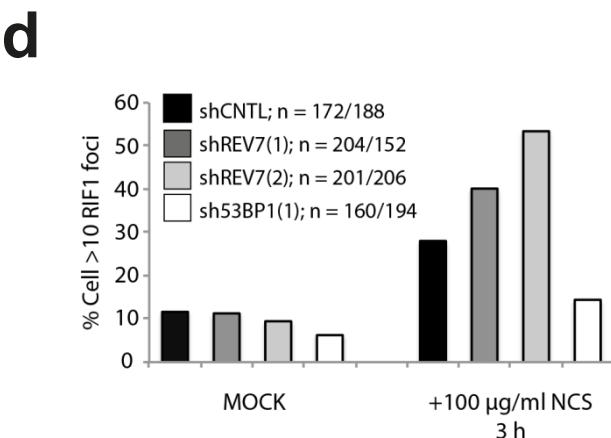
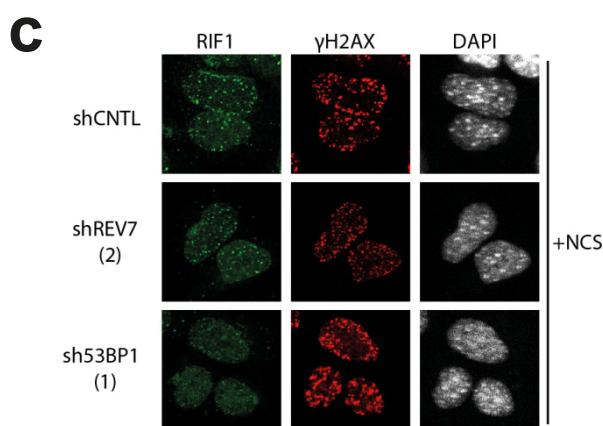
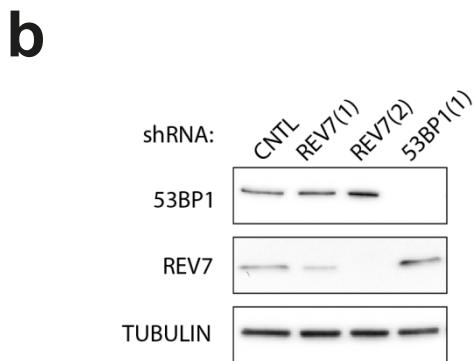
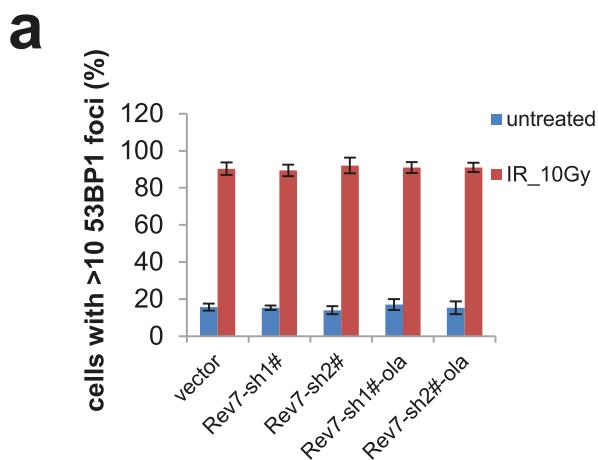
Extended Data Figure 6 | Rev7 loss promotes end resection at DSBs in BRCA1-deficient cells after ionizing radiation. **a**, Quantification of RPA-positive α tracks in KB1P-B11 cells 1 or 2 h after irradiation with α particles. **b**, Quantification of RPA- and 53BP1-positive α tracks in KB1P-B11 cells transfected with non-targeting control siRNAs or siRNAs against *Ctip*. **c**, Cell cycle analysis (BrdU incorporation and propidium iodine labelling) of KB1P-B11 cells transduced with the indicated constructs and siRNAs. **d, e**, Quantification of Rev7 transcript (**d**) or protein (**e**) levels in

BRCA1-deficient mES cells transduced with *Rev7*-targeting shRNAs or the vector control. *Hprt* was used as a control for transcript expression, α -tubulin as a control for protein expression. The data represent mean \pm s.d. **f**, Representative images of surviving colonies of *Brca1*^{-/-} mES cells transduced with an empty vector control or *Rev7*-targeting shRNAs. **g**, Quantification of colony formation normalized to the vector control. **h**, Quantification of RAD51 foci in *Brca1*^{-/-} mES cells that were transduced with the indicated constructs.

a**b****c****d****e**

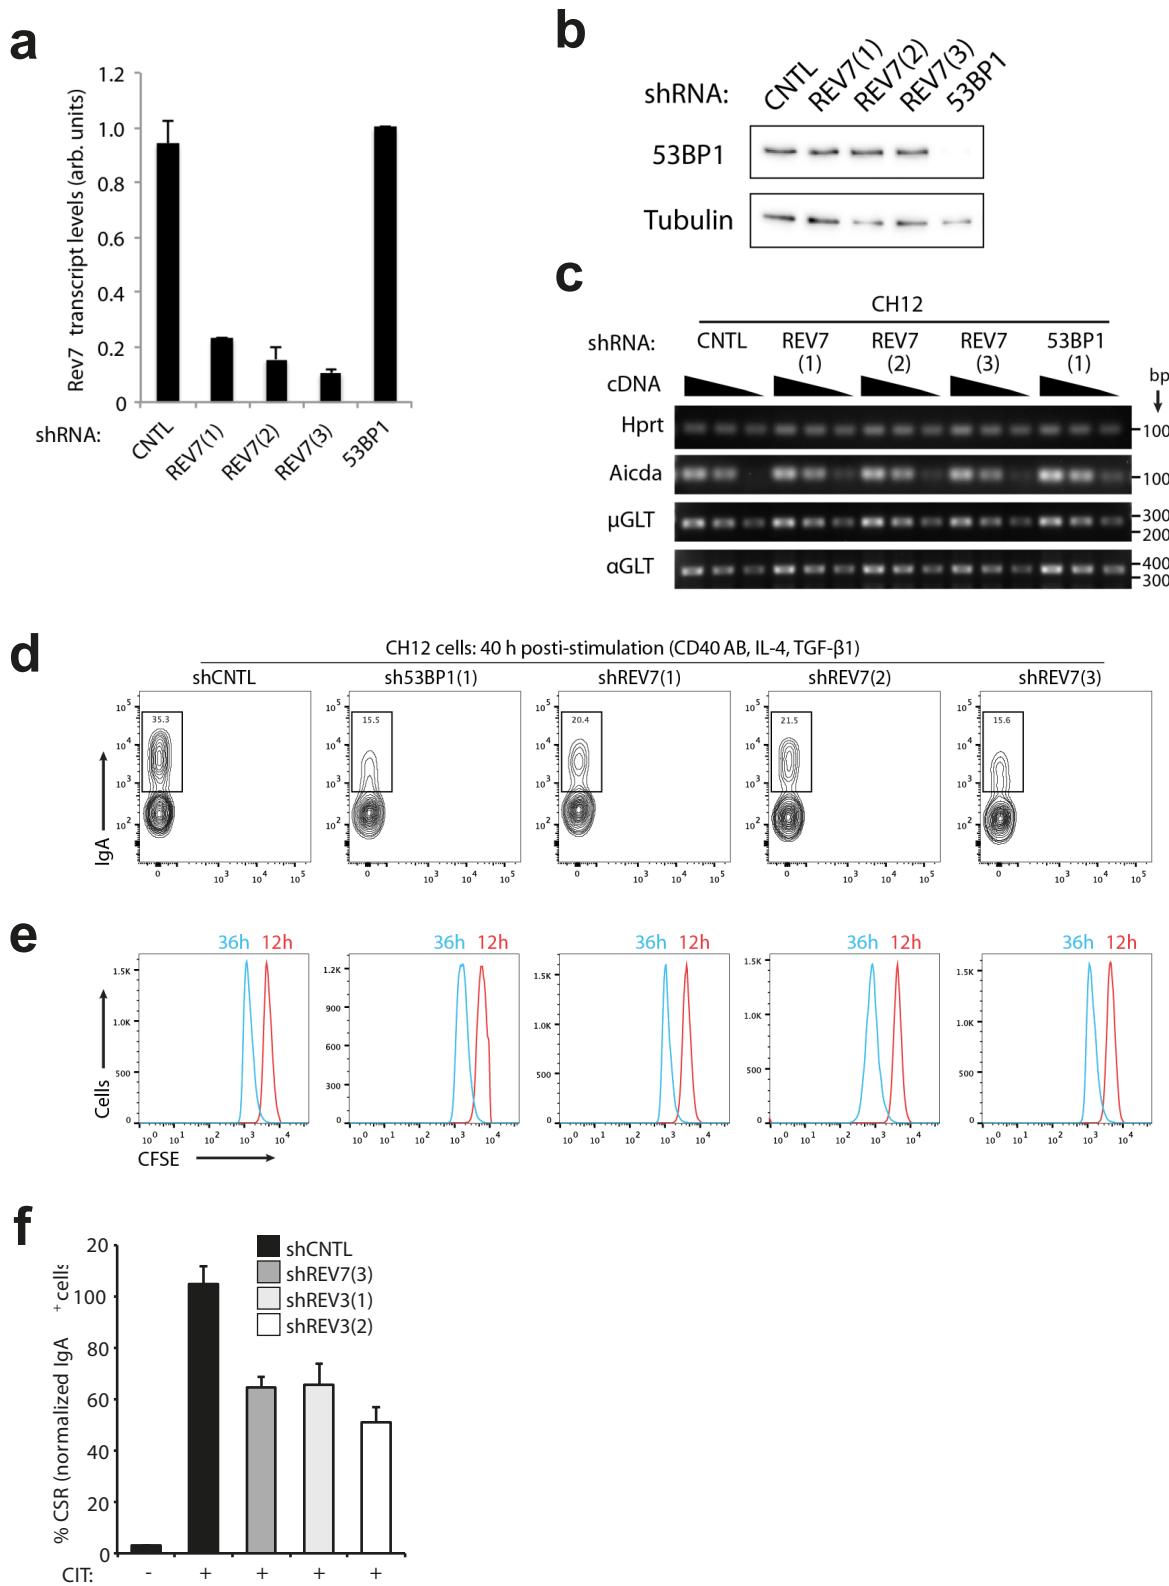
Extended Data Figure 7 | REV7 loss frequently occurs in triple-negative breast cancer. **a, b,** Quantification of human *REV7* transcript levels (**a**) and protein levels (**b**) in U2OS cells transduced with indicated constructs. Two different pairs of primers for *REV7* or *HPRT* were used for the quantification of *REV7* transcript levels. **c–e,** Examples of aberrantly reduced *REV7* protein expression in triple-negative human breast carcinomas. Immunohistochemical

detection of *REV7* in human breast carcinomas shows moderate to high nuclear expression in normal human breast tissue (data not shown), and most invasive breast tumours (**c**). Aberrant reduction of *REV7* with less than 70% of cancer cells that show nuclear positivity (**d, e**) was observed in 18 out of 50 cases.



Extended Data Figure 8 | REV7 is a downstream effector of 53BP1.
a, Quantification of 53BP1 foci in KB1P-G3 cells in the presence or absence of REV7 depletion. At least 100 cells were analysed per group in three independent experiments each. Error bars indicate s.d.; IR denotes 5 h after 10 Gy ionizing radiation. b, REV7 or 53BP1 protein levels were determined by western blotting of lysates derived from MEF cells transduced with the indicated control (CNTL) or *Rev7*- and *53bp1*-targeting shRNA constructs.

c, d, RIF1 foci formation (c) after neocarzinostatin (NCS) treatment and quantification of RIF1 foci (d) in MEF cells in the presence or absence of REV7 or 53BP1 depletion. e, Flag pull-downs were performed from 2 mg lysate prepared from *53bp1*^{-/-}, *53bp1*^{-/-} plus 53BP1 and *53bp1*^{-/-} plus 53BP1^{20AQ} MEFs after mock or neocarzinostatin treatment. Control, 53BP1 and 53BP1^{20AQ} immunoprecipitates were incubated in HeLa nuclear extract (HNE, 2 mg) and then eluted with triple-Flag peptide. HA, haemagglutinin.



Extended Data Figure 9 | The effect of REV7 inhibition on CSR after antigenic stimulation of CH12 cells. **a**, *Rev7* messenger RNA levels determined by qRT-PCR were normalized against β -actin (*Actb*) transcripts in the indicated shRNA-transduced CH12 cell lines. The data represent the mean \pm s.e.m. from two primer sets specific for *Rev7* transcript. CNTL, control. **b**, 53BP1 protein of each group normalized to vector-transduced cells (CH12) was analysed by western blotting. **c**, IgH μ and α germ-line transcripts (GLT) and *Aid* mRNA were estimated by semi-quantitative RT-PCR using twofold serial dilutions of cDNA made from indicated CH12 cell lines 40 h after

stimulation. *Hprt* was used as a control for transcript expression. **d**, Representative flow cytometric profiles of shRNA-transduced CH12 B cells stained with anti-IgA antibody 40 h after stimulation with the indicated cytokines. **e**, Cells (CH12) were labelled with CFSE immediately before cytokine stimulation as in Fig. 4d, and cell proliferation was assessed by flow cytometry at indicated time points. **f**, Quantification of CSR to IgA of shRNA-transduced CH12 cells 40 h after stimulation with CD40 antibody, IL-4 and TGF- β 1 (CIT). Data represent the mean \pm s.d. from two independent experiments performed in triplicate.

Atomic structure of anthrax protective antigen pore elucidates toxin translocation

Jiansen Jiang^{1,2}, Bradley L. Pentelute³, R. John Collier⁴ & Z. Hong Zhou^{1,2}

Anthrax toxin, comprising protective antigen, lethal factor, and oedema factor, is the major virulence factor of *Bacillus anthracis*, an agent that causes high mortality in humans and animals. Protective antigen forms oligomeric pre pores that undergo conversion to membrane-spanning pores by endosomal acidification, and these pores translocate the enzymes lethal factor and oedema factor into the cytosol of target cells¹. Protective antigen is not only a vaccine component and therapeutic target for anthrax infections but also an excellent model system for understanding the mechanism of protein translocation. On the basis of biochemical and electrophysiological results, researchers have proposed that a phi (Φ)-clamp composed of phenylalanine (Phe)427 residues of protective antigen catalyses protein translocation via a charge-state-dependent Brownian ratchet^{2–9}. Although atomic structures of protective antigen pre pores are available^{10–14}, how protective antigen senses low pH, converts to active pore, and translocates lethal factor and oedema factor are not well defined without an atomic model of its pore. Here, by cryo-electron microscopy with direct electron counting, we determine the protective antigen pore structure at 2.9-Å resolution. The structure reveals the long-sought-after catalytic Φ -clamp and the membrane-spanning translocation channel, and supports the Brownian ratchet model for protein translocation. Comparisons of four structures reveal conformational changes in pre pore to pore conversion that support a multi-step mechanism by which low pH is sensed and the membrane-spanning channel is formed.

Triggering conversion from protective antigen (PA) pre pore to pore by *in vitro* acidification leads to rapid and irreversible aggregation. Attempts to prevent aggregation by screening detergents have largely failed¹⁵. By low-pH treatment of PA pre pores directly on electron microscopy (EM) grids containing a thin layer of continuous carbon film, we obtained dispersed particles of PA pore without aggregation (Extended Data Fig. 1). We then acquired drift-corrected cryo-electron microscopy (cryoEM) images (Extended Data Fig. 1b–d) and reconstructed a map at an overall resolution of 2.9 Å using 60,455 particles (Fig. 1, Extended Data Fig. 2 and Supplementary Video 1). The resolution for most regions of the cryoEM map is ~2.8 Å (Extended Data Fig. 2c). Our map reveals rich high-resolution structural features, including amino-acid side chains and 14 chelated Ca²⁺ ions (Extended Data Fig. 1e–h), and has allowed unambiguous *de novo* atomic modelling (Extended Data Table 1) and detailed structure and function analyses.

The overall structure of the PA pore has a ‘flower-on-a-stem’ architecture, including corolla, calyx, and stem from top to bottom (Figs 1b and 2a and Supplementary Video 1). Each PA protomer is divided into four domains in the PA pre pore¹⁰, named 1', 2, 3, and 4. In the PA pore, domains 1', 3, and 4 form the corolla and domain 2 forms the calyx and the stem; therefore we designate the parts of domain 2 corresponding to the calyx and the stem as 2c (residues 259–274 and 354–487) and 2s (residues 275–353), respectively (Fig. 2b, c). Domains 1' and 2c form a compact structure responsible for substrate protein binding and intake (Fig. 2). Domain 2s is an extended β -hairpin (2 β _{2s} and 2 β _{3s}), seven

copies of which assemble to form a membrane-spanning 14-stranded β -barrel 105 Å in length and 27 Å (from C_α to C_α) in diameter (Fig. 2). Domain 3 is located peripherally and has close contact with domains 1' and 2c (Fig. 2b). The cryoEM density of domain 4 is weak and has the lowest resolution among all domains (inset of Fig. 1b and Extended Data Fig. 2c), probably because of its flexibility from minimal contact with the other domains. Rigid-body fitting of domain 4 of the PA pre pore crystal structure to the cryoEM map shows domain 4 shifts ~4 Å towards the central axis in the PA pore (Extended Data Fig. 3).

The translocation channel of the PA pore has a funnel shape and can be divided into four parts based on diameters calculated with MOLE¹⁶: mouth, Φ -clamp, throat, and tube (Fig. 3a, b). Its surface is negatively charged and mainly hydrophilic, but hydrophobic patches are seen at the α -clamp¹⁴ of the mouth, near the Φ -clamp, and at the middle of the tube (Fig. 3a, c). As proposed on the basis of the PA pre pore structure¹⁰, the negatively charged surface promotes passage of cations, and the hydrophilic surface facilitates passage of substrate proteins/polypeptides.

The mouth has a 30-Å opening and inner diameters varying down to 20 Å (Fig. 3a); it can accommodate protein secondary structure elements, but not folded domains such as the amino (N)-terminal of lethal factor (LF_N). The Φ -clamp below the mouth becomes the bottleneck of the entire channel, with a solvent-excluded inner diameter of only 6 Å (Fig. 3a, b), which is smaller than protein secondary structure elements and therefore may only allow passage of fully unfolded polypeptides. Underneath the Φ -clamp are the throat, which is an enlarged (~18 Å) bulb-shaped chamber, and the tube formed by the 14-stranded β -barrel with inner diameters in the range 12–18 Å and rich in Ser and Thr residues in its middle region (Fig. 3a, b). The large diameter (>12 Å) of

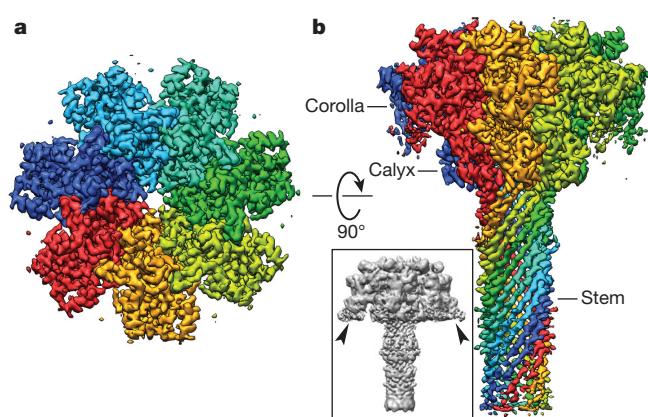


Figure 1 | CryoEM reconstruction of the PA pore. **a, b,** Surface representations of the cryoEM map of the PA pore at 2.9-Å resolution as viewed from the top (a) and the side (b). Individual protomers of PA heptamer are colour-coded. Inset of b shows the unsharpened cryoEM map in which the flexible domains 4 (arrowheads) are visible.

¹Department of Microbiology, Immunology and Molecular Genetics, University of California, Los Angeles, California 90095, USA. ²California NanoSystems Institute, University of California, Los Angeles, California 90095, USA. ³Department of Chemistry, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA. ⁴Department of Microbiology and Immunobiology, Harvard Medical School, Boston, Massachusetts 02115, USA.

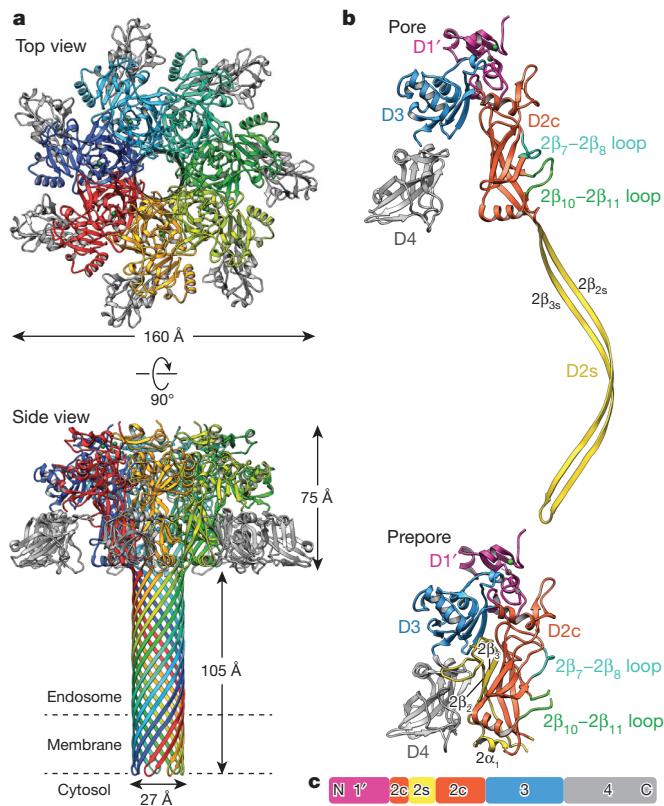


Figure 2 | Atomic model of the PA pore. **a**, Top and side views of the atomic model of the PA pore shown as ribbons. Protomers are colour-coded except for domain 4 (grey). **b**, Structural comparison of the protomers of the PA pore and prepore (PDB accession number 1TZO). The domains are coloured differently according to **c**. **c**, Domain organization of the PA protomer.

the throat and tube can accommodate an α -helix which may be formed by polypeptides after passing the Φ -clamp (Fig. 3a). The diameter and the hydrophilic property of this part of channel are similar to those of the exit tunnel of the ribosome, which is proposed to translocate α -helices^{17,18}.

The mouth and the tube are the only two openings of the channel, and are accessible to the endosomal and cytosolic compartments, respectively (Fig. 3a). The rest of the channel is ‘water-tight’ and without holes permeable to small molecules. The substrate protein blocks the small hole of the Φ -clamp before or during translocation². Thus the Φ -clamp may act as a gate separating the endosomal and cytosolic compartments. The differences of proton concentration ($\Delta p\text{H}$) and electrical potential ($\Delta\psi$) across the Φ -clamp (Fig. 3a) may serve as the primary driving force for substrate protein translocation³.

In contrast to the hydrophilic inner surface of the β -barrel, its outer surface is largely hydrophobic (Fig. 3c, d). The Phe residues (313, 314, and 324) form two aromatic belts on opposite sides of the lipid bilayer (Fig. 3d), which may stabilize membrane insertion^{19,20}. The cryoEM map shows that the transmembrane region is surrounded by a cloud of disordered densities, which we interpret as bound detergent molecules that were added during sample preparation (Fig. 3e). Surprisingly, an additional hydrophobic surface in the middle of the β -barrel containing a cluster of hydrophobic residues (Ile289, Val332, Ile334, Leu338 and Leu340 from each protomer) was also bound with detergent molecules (Fig. 3d–f). However, the function of this additional hydrophobic region is unknown.

The overall architecture of the PA pore is similar to those of bacterial toxins α -haemolysin and *Vibrio cholerae* cytolsin in their membrane-inserted conformations (Extended Data Fig. 4) despite differences in functions^{21,22}. The β -barrel of the PA pore is twice as long as those of the other two toxins and may facilitate LF and oedema factor translocation by entropically stabilizing an α -helix within the confinement of

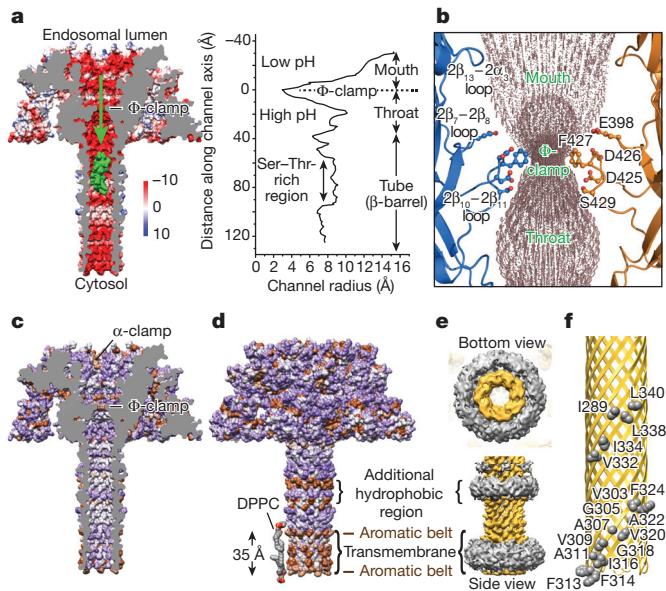


Figure 3 | Translocation channel of the PA pore. **a**, Electrostatic surface (left) and channel radius profile (right) of the PA pore. An α -helix (green surface) model of residues 555–574 of lethal factor; PDB accession number 1J7N is modelled in the β -barrel. The green arrow depicts the direction of protein translocation. Colour scale, kcal (mol e)⁻¹. **b**, The translocation channel (dots) running through the Φ -clamp. Two protomers of the PA pore are shown as ribbons. Residues Glu398, Asp425, Asp426, Phe427, and Ser429 (ball-and-stick) are exposed to the channel. **c**, **d**, Hydropathy surface of the PA pore (brown: hydrophobic; purple: hydrophilic; white: neutral). In **c**, the front half of the structure is removed to show the luminal surface. In **d**, two 1,2-dihexadecanoyl-sn-glycero-3-phosphocholine (DPPC) lipid molecules are modelled near the membrane insertion region. **e**, Bottom and side views of the segmented cryoEM map showing the β -barrel (yellow) bound with disordered detergent molecules (grey). **f**, The 14-stranded β -barrel (ribbons) and the hydrophobic residues (spheres) on its outer surface. The hydrophobic residues are depicted on different protomers for the ease of presentation.

the cylindrical channel^{9,17}. This longer barrel may also be necessary to accommodate its receptor situated between domain 4 and the host membrane¹¹ (Extended Data Fig. 4). Except for this length difference, the 14-stranded β -barrels of the PA pore, α -haemolysin, and *V. cholerae* cytolsin share geometries, such as diameter, twist of β -strands, and pitch length (Extended Data Fig. 4).

In the PA pore, the seven 2 β ₁₀–2 β ₁₁ loops of the heptamer converge to form an iris with a 6-Å hole bounded by a symmetrical arrangement of the seven Phe427 residues (Fig. 4). By contrast, these loops in the PA prepore do not engage each other^{10,12}, but circumscribe a 30-Å central hole with adjacent Phe427 residues spaced ~14 Å apart (Extended Data Fig. 5a). In the PA pore, each 2 β ₁₀–2 β ₁₁ loop is stabilized by close interactions with the 2 β ₇–2 β ₈ loop of the same protomer and the 2 β ₁₀–2 β ₁₁ loops of its two neighbouring protomers. The hydrogen bonds between Asn399 and Ser428 and between Asn399 and Lys397' of an adjacent protomer form a chain of interactions, which give rise to a ring of loops at the iris inside the PA pore (Extended Data Fig. 5b). Consistently, mutation of Ser428 abolished protein translocation, and mutation of Lys397 or Asn399 resulted in a dominant negative effect^{23–25}. In the PA homologues found in *Clostridium* species, Lys397 and Asp426 are both replaced by uncharged Gln²⁶, implying that this chain of interactions might be formed differently in these homologues.

The aryl plane of Phe427 from each protomer is parallel to the central axis of the PA pore. Neighbouring Phe427 residues interact with each other by aromatic CH–π interaction in a tilted T-shaped configuration²⁷ and possibly by hydrophobic interaction, thus forming the Φ -clamp (Fig. 4 and Extended Data Fig. 5). The integrity of the Φ -clamp is required for catalysis of protein translocation, as mutation of even a single Phe427

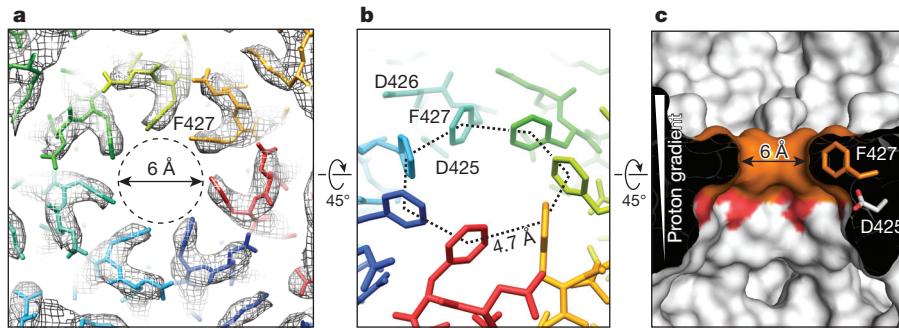


Figure 4 | Structure of the Φ -clamp. **a**, Top view of the Φ -clamp region of the PA pore showing the cryoEM map (mesh) superimposed with the atomic model (stick). **b**, Tilted view of the Φ -clamp with seven protomers coloured differently, showing the aromatic CH- π interaction. **c**, Cross-section side view

of the translocation channel near the Φ -clamp region. The Φ -clamp (Phe427) and the conserved acidic residue Asp425 are coloured in orange and red, respectively.

residue of the Φ -clamp severely reduced translocation efficiency and disrupted the seal against cation passage⁵. Because the PA pore translocates polypeptides with various side chains, the Φ -clamp may act like an elastic ‘O-ring’, changing its size and shape to allow an unhindered passage of different amino-acid residues while maintaining a good seal during translocation. In the recently reported structure of bacterial amyloid secretion channel CsgG, eight Phe residues form a Φ -clamp, which differs from that in the PA pore in having a 9.5-Å hole and facilitating protein translocation in an ungated manner²⁸.

Evidence from electrophysiological studies indicates that the proton gradient across the endosomal membrane is the primary driving force for unidirectional translocation of proteins through the PA pore³. Proton-driven transporters usually involve protonation-dependent conformational changes of two or more alternating gates. In the PA pore, however, there is only one gate, namely the Φ -clamp, and neither the Φ -clamp nor its nearby residues probably undergo a protonation-dependent conformational change. Indeed, the PA pore structure supports the charge-state-dependent Brownian ratchet model^{2–9}, proposed earlier. In this model, a negative electrostatic barrier within the pore hinders the passage, by Brownian motion, of deprotonated acidic residues^{3,4,29}. The fact that acidic residues in the acidic environment

of the endosome have a higher probability of being protonated, and are thus free to pass the barrier, than those in the neutral environment of the cytosol, leads necessarily to unidirectional movement of polypeptides across the barrier. Consistent with this model, the PA pore structure shows three acidic residues, Asp425, Asp426, and Glu398, to be near the Φ -clamp and proximal to the pore axis (Fig. 3b), generating a strong negative electrostatic barrier demarcating the endosomal and cytosolic compartments. In addition, the highly conserved acidic residue Asp425 located directly underneath the Φ -clamp is ideally positioned as a proton sink that may strip off protons from protonated acidic residues passing the Φ -clamp (Fig. 4c). Besides the charge barrier in the Φ -clamp, efficient protein translocation may require additional charged spots, such as the top region of the β -barrel⁸.

A key question about the low-pH-triggered conversion from prepose to pore is how PA senses pH. A notable conformational change during the conversion is that the $2\beta_{10}$ - $2\beta_{11}$ loop is flipped from one side to the other (Fig. 5a and Extended Data Fig. 6). Interestingly, this loop has different conformations in the two crystal structures of PA octameric preposes ([PA₆₃]₈, PDB accession number 3HVD; [PA₆₃]₈[LF_N]₄, PDB accession number 3KWF) whereas the remaining parts of the structures are largely unchanged^{9,13,14}; that is, this loop is in the pore and

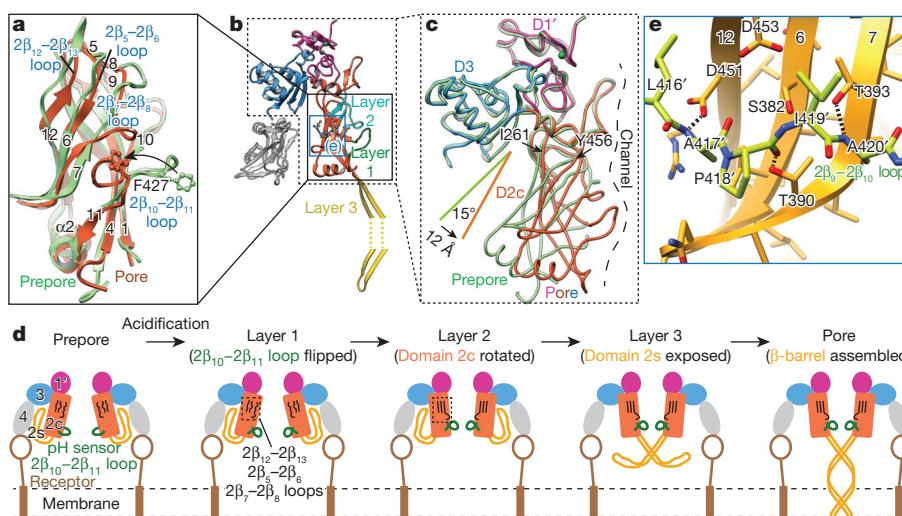


Figure 5 | Conversion from prepose to pore. **a**, Superimposition of domains 2c of the PA pore and prepose (monomer; PDB accession number 3TEW). Loops undergoing dramatic conformational changes during the prepose to pore conversion are labelled (blue). **b**, Overview of the three layers of conformational changes during the conversion. The interaction between the $2\beta_9$ - $2\beta_{10}$ loop and the β -sheet (blue box) is shown in **e**. **c**, Superimposition of the PA pore (pink, orange, and blue; coloured by domain) and prepose (light green; PDB accession number 1TZQ). Domains 4 and 2s are removed for

clarity. Between the two structures, domain 2c has a $C\alpha$ root-mean-square deviation = 6.87 Å, whereas domains 1' and 3 remain unchanged with a $C\alpha$ root-mean-square deviation = 0.54 Å and 0.58 Å, respectively. **d**, Illustration of the steps of conformational changes during the prepose to pore conversion. For simplicity, only two protomers are illustrated. **e**, Interface between the $2\beta_9$ - $2\beta_{10}$ loop (green) and the bent β -sheet (gold) of $2\beta_{12}$, $2\beta_6$, and $2\beta_7$ of the adjacent protomer in the PA pore. Hydrogen bonds between the backbone of the loop and the side chains of the β -sheet are depicted with dashed lines.

prepore states in 3HVD (Extended Data Fig. 6b) and 3KVV, respectively. This comparison suggests that the $2\beta_{10}$ - $2\beta_{11}$ loop can switch between two conformations without affecting the overall structure. Mutation of the conserved Asn422 or Asp425 in this loop abolished the prepore to pore conversion at low pH in a dominant negative manner²⁵. Therefore we interpret this loop as a pH sensor. Surprisingly, this loop does not possess any His, a residue whose pK_a of approximately 6 falls in the pH range of 5–7 where the prepore to pore conversion occurs. The conformational change of this loop might result from destabilization and rearrangement of hydrogen bonds upon exposure to low pH, or alternatively from perturbed pK_a of Asp residues (425 and 426) in this loop. We note that other parts of PA might impact pH sensing. Prior work has shown mutation of residues near or in the $2\beta_2$ - $2\beta_3$ loop, or at the domain 2/domain 4 interface, or deletion of the $2\beta_2$ - $2\beta_3$ loop changed the pH threshold for conversion but did not abolish the low-pH sensitivity^{25,30}. These observations suggest that the involved residues might regulate pH sensing.

We find that the prepore to pore conversion involves three layers of conformational changes, all located within domain 2 (Fig. 5b). The first layer is the above-mentioned pH-sensing $2\beta_{10}$ - $2\beta_{11}$ loop. Low-pH-triggered conformational change of this loop is relayed to the following layers sequentially as detailed below.

The second layer involves the $2\beta_7$ - $2\beta_8$, $2\beta_5$ - $2\beta_6$, and $2\beta_{12}$ - $2\beta_{13}$ loops, which are positioned in an anti-parallel manner (Fig. 5a). In the PA pore, the N-terminal half of the $2\beta_7$ - $2\beta_8$ loop moves downwards to form a U-shaped turn tethered to the putative pH-sensing $2\beta_{10}$ - $2\beta_{11}$ loop via a hydrogen bond between Asn399 and Ser428 (Fig. 5a and Extended Data Figs 5b and 6e). This new inter-loop interaction may relay the low-pH-triggered conformational changes from the first layer to the second layer as supported by the dominant negative effect of Asn399 or Lys397 mutation^{24,25}, which breaks the hydrogen bond network connecting these two loops (Extended Data Fig. 5b). As a result of the downward pull from the $2\beta_7$ - $2\beta_8$ loop, all three loops of the second layer become straightened in the PA pore, with parts of the $2\beta_5$ - $2\beta_6$ and $2\beta_{12}$ - $2\beta_{13}$ loops forming strands to augment $2\beta_6$ and $2\beta_{12}$, respectively (Fig. 5a and Extended Data Fig. 6c, f). This conformational change is supported by the observation that mutating any of Pro379, Val455, and Asn458 of these loops led to defects in protein translocation²⁵. This conformational change leads to the convergence of domain 2c of each protomer and formation of the Φ -clamp by a 15° rotation of domain 2c around the hinge residues Ile261 and Tyr456, with its distal edge moving for 12 Å towards the central axis of the PA pore (Fig. 5c and Supplementary Video 2). Indeed, tethering domain 2 to domain 4 by receptor binding or stabilizing the interface between these two domains by a poly- γ -D-glutamate capsule impeded domain convergence and therefore decreased the pH threshold required for the prepore to pore conversion^{11,12,30}. Conversely, destabilizing the domain 2/domain 4 interface by point mutations increased the pH threshold³⁰.

The third layer, resulting from the second layer of conformational changes, involves release and refolding of the precursor of domain 2s (that is, $2\beta_2$, $2\beta_3$, $2\alpha_1$, and the membrane insertion loop), and leads to the formation of the membrane-spanning β -barrel (Figs 2 and 5b and Extended Data Fig. 7). The rotation of domain 2c opens the pocket between domain 2c and domain 4 to release the precursors of domain 2s, which then come together to form the β -barrel (see details in Extended Data Fig. 7b). The β -barrel formation might be the result, rather than the cause, of the convergence of domain 2c, as supported by a cryoEM map obtained from a subset of particles that shows converged domain 2c in the pore state but lacks the β -barrel (Extended Data Fig. 8). We note that the stability of the domain 2/domain 4 interface was found to be a rate-limiting barrier to the conversion³⁰ and the separation between domain 2c and domain 4 might also happen at an earlier stage of the conversion.

Taken together, these results suggest a multi-step mechanistic model of the low-pH-triggered conversion from prepore to pore (Fig. 5d). First, the $2\beta_{10}$ - $2\beta_{11}$ loop as a pH sensor changes its conformation upon

acidification (layer 1). Second, consequently, the $2\beta_7$ - $2\beta_8$, $2\beta_5$ - $2\beta_6$, and $2\beta_{12}$ - $2\beta_{13}$ loops become ordered (layer 2), resulting in the convergence of domain 2c and the formation of the Φ -clamp. Last, as a result of the separation of domain 2c from domain 4, the precursors of domain 2s are released and refold into a uniform β -barrel (layer 3), which, as a final step, inserts into the endosomal membrane, ready for delivering the toxic enzymes into the cytosol.

In the PA pore, the convergence of domain 2c creates a large interprotomer interface to stabilize the pore conformation (Extended Data Fig. 7a). Key to the formation of a functional PA pore is a newly identified interface between the $2\beta_9$ - $2\beta_{10}$ loop of one protomer and the triple-stranded β -sheet ($2\beta_{12}$, $2\beta_6$, and $2\beta_7$) of its neighbour (Fig. 5e). The above-mentioned convergence of domain 2c brings them together to form a new binding interface primarily mediated by hydrogen bonds between the backbone of the loop and the side chains of Thr390, Thr393, and Asp451 of the β -sheet (Fig. 5e). These interactions explain the dominant negative effect of Ser382 or Thr393 mutant, the loss of PA activity of the Cys mutant of Thr390 or Asp451, and the conservation of Thr390 and Thr393 in PA homologues²⁵. By contrast, these residues are accessible and free of interactions in the PA prepore, and thus might be potential drug targets for blocking the conversion to functional PA pores.

In summary, the 2.9-Å structure of the anthrax PA pore suggests the $2\beta_{10}$ - $2\beta_{11}$ loop as a pH sensor to trigger conformational changes for prepore to pore conversion, supports a charge-state-dependent Brownian ratchet model of proton-driven protein translocation, and can inform efforts both to engineer PA to target cancer cells and to design measures to block anthrax toxin entering cells^{4,21}. A more detailed mechanism of protein translocation through the PA pore awaits additional experimental evidence, such as structures of the PA pore in the act of polypeptide translocation.

Online Content Methods, along with any additional Extended Data display items and Source Data, are available in the online version of the paper; references unique to these sections appear only in the online paper.

Received 28 October 2014; accepted 19 January 2015.

Published online 16 March 2015.

- Young, J. A. & Collier, R. J. Anthrax toxin: receptor binding, internalization, pore formation, and translocation. *Annu. Rev. Biochem.* **76**, 243–265 (2007).
- Krantz, B. A. *et al.* A phenylalanine clamp catalyzes protein translocation through the anthrax toxin pore. *Science* **309**, 777–781 (2005).
- Krantz, B. A., Finkelstein, A. & Collier, R. J. Protein translocation through the anthrax toxin transmembrane pore is driven by a proton gradient. *J. Mol. Biol.* **355**, 968–979 (2006).
- Basilio, D., Juris, S. J., Collier, R. J. & Finkelstein, A. Evidence for a proton-protein symport mechanism in the anthrax toxin channel. *J. Gen. Physiol.* **133**, 307–314 (2009).
- Janowik, B. E., Fischer, A. & Collier, R. J. Effects of introducing a single charged residue into the phenylalanine clamp of multimeric anthrax protective antigen. *J. Biol. Chem.* **285**, 8130–8137 (2010).
- Basilio, D., Jennings-Antipov, L. D., Jakes, K. S. & Finkelstein, A. Trapping a translocating protein within the anthrax toxin channel: implications for the secondary structure of permeating proteins. *J. Gen. Physiol.* **137**, 343–356 (2011).
- Brown, M. J., Thoren, K. L. & Krantz, B. A. Charge requirements for proton gradient-driven translocation of anthrax toxin. *J. Biol. Chem.* **286**, 23189–23199 (2011).
- Wynia-Smith, S. L., Brown, M. J., Chirichella, G., Kemalyan, G. & Krantz, B. A. Electrostatic ratchet in the protective antigen channel promotes anthrax toxin translocation. *J. Biol. Chem.* **287**, 43753–43764 (2012).
- Feld, G. K., Brown, M. J. & Krantz, B. A. Ratcheting up protein translocation with anthrax toxin. *Protein Sci.* **21**, 606–624 (2012).
- Petosa, C., Collier, R. J., Klimpel, K. R., Leppla, S. H. & Liddington, R. C. Crystal structure of the anthrax toxin protective antigen. *Nature* **385**, 833–838 (1997).
- Santelli, E., Bankston, L. A., Leppla, S. H. & Liddington, R. C. Crystal structure of a complex between anthrax toxin and its host cell receptor. *Nature* **430**, 905–908 (2004).
- Lacy, D. B., Wigelsworth, D. J., Melnyk, R. A., Harrison, S. C. & Collier, R. J. Structure of heptameric protective antigen bound to an anthrax toxin receptor: a role for receptor in pH-dependent pore formation. *Proc. Natl. Acad. Sci. USA* **101**, 13147–13151 (2004).
- Kintzer, A. F. *et al.* The protective antigen component of anthrax toxin forms functional octameric complexes. *J. Mol. Biol.* **392**, 614–629 (2009).
- Feld, G. K. *et al.* Structural basis for the unfolding of anthrax lethal factor by protective antigen oligomers. *Nature Struct. Mol. Biol.* **17**, 1383–1390 (2010).

15. Vernier, G. et al. Solubilization and characterization of the anthrax toxin pore in detergent micelles. *Protein Sci.* **18**, 1882–1895 (2009).
16. Sehnal, D. et al. MOLE 2.0: advanced approach for analysis of biomacromolecular channels. *J Cheminform.* **5**, 39 (2013).
17. Ziv, G., Haran, G. & Thirumalai, D. Ribosome exit tunnel can entropically stabilize alpha-helices. *Proc. Natl Acad. Sci. USA* **102**, 18956–18961 (2005).
18. Voss, N. R., Gerstein, M., Steitz, T. A. & Moore, P. B. The geometry of the ribosomal polypeptide exit tunnel. *J. Mol. Biol.* **360**, 893–906 (2006).
19. Kilian, J. A. & von Heijne, G. How proteins adapt to a membrane-water interface. *Trends Biochem. Sci.* **25**, 429–434 (2000).
20. Wang, J., Vernier, G., Fischer, A. & Collier, R. J. Functions of phenylalanine residues within the beta-barrel stem of the anthrax toxin pore. *PLoS ONE* **4**, e6280 (2009).
21. Song, L. et al. Structure of staphylococcal α -hemolysin, a heptameric transmembrane pore. *Science* **274**, 1859–1866 (1996).
22. De, S. & Olson, R. Crystal structure of the *Vibrio cholerae* cytolytic heptamer reveals common features among disparate pore-forming toxins. *Proc. Natl Acad. Sci. USA* **108**, 7385–7390 (2011).
23. Sellman, B. R., Nassi, S. & Collier, R. J. Point mutations in anthrax protective antigen that block translocation. *J. Biol. Chem.* **276**, 8371–8376 (2001).
24. Sellman, B. R., Mourez, M. & Collier, R. J. Dominant-negative mutants of a toxin subunit: an approach to therapy of anthrax. *Science* **292**, 695–697 (2001).
25. Mourez, M. et al. Mapping dominant-negative mutations of anthrax protective antigen by scanning mutagenesis. *Proc. Natl Acad. Sci. USA* **100**, 13803–13808 (2003).
26. Melnyk, R. A. & Collier, R. J. A loop network within the anthrax toxin pore positions the phenylalanine clamp in an active conformation. *Proc. Natl Acad. Sci. USA* **103**, 9802–9807 (2006).
27. Ballester, P. & Biros, S. M. in *The Importance of Pi-Interactions in Crystal Engineering* 79–107 (John Wiley, 2012).
28. Goyal, P. et al. Structural and mechanistic insights into the bacterial amyloid secretion channel CsgG. *Nature* **516**, 250–253 (2014).
29. Finkelstein, A. Proton-coupled protein transport through the anthrax toxin channel. *Phil. Trans. R. Soc. B* **364**, 209–215 (2009).
30. Kintzer, A. F., Tang, H., Schawel, A. K., Brown, M. J. & Krantz, B. A. Anthrax toxin protective antigen integrates poly-gamma-D-glutamate and pH signals to sense the optimal environment for channel formation. *Proc. Natl Acad. Sci. USA* **109**, 18378–18383 (2012).

Supplementary Information is available in the online version of the paper.

Acknowledgements This work was supported in part by grants from the National Institutes of Health (AI094386/AI046420 and GM071940 to Z.H.Z., and AI022021 to R.J.C.), the American Heart Association (Postdoctoral Fellowship 14POST18870059 to J.J.), Damon Runyon Cancer Research Foundation Innovation Award (B.L.P.), and National Science Foundation (CAREER Award CHE-1351807 to B.L.P.). We thank the NERCE facility (the National Institutes of Health grant AI057159) for expression of toxin proteins. We acknowledge the use of instruments at the Electron Imaging Center for Nanomachines supported by University of California, Los Angeles, and by instrumentation grants from the National Institutes of Health (1S10RR23057, 1S10OD018111) and National Science Foundation (DBI-1338135). We thank L. Jin for initial efforts on this project and J. Feigon for discussion.

Author Contributions J.J. designed and performed the experiments, analysed data, and wrote the paper; Z.H.Z. initialized and supervised the research, analysed data, and wrote the paper. B.L.P. and R.J.C. analysed data and wrote the paper.

Author Information Three-dimensional cryoEM density maps of anthrax PA pore have been deposited in the Electron Microscopy Data Bank under accession numbers EMD-6224 (intact PA pore) and EMD-6225 (lacking the membrane-spanning β -barrel). The coordinates of atomic model of anthrax PA pore have been deposited in the Protein Data Bank under accession number 3J9C. Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to Z.H.Z. (Hong.Zhou@ucla.edu).

METHODS

No statistical methods were used to predetermine sample size.

Preparation of PA precores. PA₆₃ precores were prepared following the described procedures³¹.

EM sample preparation and data acquisition. For negative-stain EM, 2 µl of 0.1% polylysine solution (Polysciences) was first applied to a glow-discharged grid covered with carbon film and then removed by blotting with filter paper in 2 min. The polylysine treatment produced different orientations of particles on EM grids³², therefore overcoming the problem of preferred orientation for the PA pore. Immediately after removal of polylysine, 2 µl of PA prepore (~50 µg ml⁻¹) was applied to the grid and incubated for 1 min. The grid was then washed with the high-pH buffer (50 mM HEPES, 50 mM NaCl, pH 8.0) twice followed by two washes with the low-pH buffer (50 mM NaOAc, pH 5.0, 0.05% Igepal CA-630) to induce the conversion of prepore to pore. After removal of excess buffer, the grid was stained with 0.8% (w/v) uranyl formate.

For cryoEM, Quantifoil R1.2/1.3 holey grids were covered with a thin layer of continuous carbon film one day before use. The procedure for pore induction on cryoEM grids followed the same procedure for the negative-stain EM except that the last step of staining was not used. About 1.5 µl of low-pH buffer was left on the grids before they were transferred into an FEI Vitrobot Mark IV. The grids were then blotted and flash-frozen in liquid ethane in the Vitrobot at 100% humidity. The frozen grids were stored in liquid nitrogen before use.

Negative-stain EM micrographs were acquired with Leginon automation software^{33,34} and a TIETZ F415MP 16-megapixel CCD camera at ×68,027 magnification in an FEI Tecnai F20 electron microscope operated at 200 kV. The micrographs were saved with 2× binning to yield a pixel size of 4.4 Å.

Frozen-hydrated cryoEM grids were loaded into an FEI Titan Krios electron microscope operated at 300 kV for automated image acquisition with Leginon. CryoEM micrographs were recorded on a Gatan K2 Summit direct electron detection camera using the electron counting mode at ×22,500 nominal magnification (calibrated pixel size of 1.28 Å on the sample level) and defocus values ranging from -1.8 to -5.1 µm. The dose rate on the camera was set to about 8 electrons per pixel per second. The total exposure time was 8 s and fractionated into 32 frames of subimages with 0.25 s exposure time for each frame. Frame images were aligned and averaged for correction of beam-induced drift using the GPU-accelerated program from Y. Cheng's laboratory³⁵. The average images from all frames were used for defocus determination and particle picking, and those from the first 16 frames (corresponding to about 20 electrons per square ångström total dose on sample) were used for two- and three-dimensional image classifications. In total, 12,416 micrographs were taken in a continuous session. The best 7,062 micrographs were selected for the following in-depth data processing.

Image processing. For negative-stain EM single particle reconstruction, 140,775 particles were picked from 1,115 negative-stain EM micrographs using the batchboxer program of EMAN³⁶. Particles were windowed out in 96 pixels × 96 pixels. The defocus value of each micrograph was determined by CTFFIND³⁷ and particles were corrected for contrast transfer function (CTF) by phase-flipping with the corresponding defocus and astigmatism values using Bsoft³⁸. An initial model was generated using the startcsym program of EMAN. The refinement was then performed with sevenfold symmetry using EMAN.

For cryoEM single particle reconstruction, 21,200 particles (256 pixels × 256 pixels) were initially picked by hand from 1,928 micrographs and subjected to auto-refinement by RELION^{39,40} using the negative-stain EM map obtained above as the initial model. The resulting cryoEM map (approximately 4-Å resolution) was used to generate projections that were then served as templates to pick 259,719 particles from all of the 7,062 micrographs using the batchboxer of EMAN. In our procedure, the defocus values of the micrographs were determined by CTFFIND and particles were corrected for CTF by phase-flipping using Bsoft. The particles were processed with two- and three-dimensional classifications using the recommended procedures of RELION (http://www2.mrc-lmb.cam.ac.uk/relion/index.php/Recommended_procedures). Two-dimensional class averages and three-dimensional class reconstructions were inspected and those without high-resolution and interpretable features were considered as 'bad' classes. Particles contributing to the bad classes were discarded. The remaining 60,455 particles were selected for the final structure refinement. The C7 symmetry was applied throughout the three-dimensional classification and three-dimensional auto-refinement.

To maximize usable signals from the frame images acquired with the K2 Summit camera, we used the resolution and dose-dependent model of radiation damage recently introduced in RELION-1.3 in the following steps⁴¹. First, the particle images averaged from all 32 frames with whole-image drift correction were used for a preliminary three-dimensional auto-refinement. Second, particle images from individual frames were used to calculate translational alignments for the particle-based drift correction. A running average of seven frames, a standard deviation of one

pixel, and fitting of linear tracks through the translations for all running averages were used for the optimal translational alignment following the suggested protocol of RELION. Last, particle images from frame 3 to frame 27 (~30 electrons per square ångström total dose on sample) were translated using the above optimal alignment and weighted with different B-factors as estimated from the single-frame reconstructions to generate optimal 'shiny' average images. Application of this procedure to the above selected particles yielded 60,455 'shiny' particles.

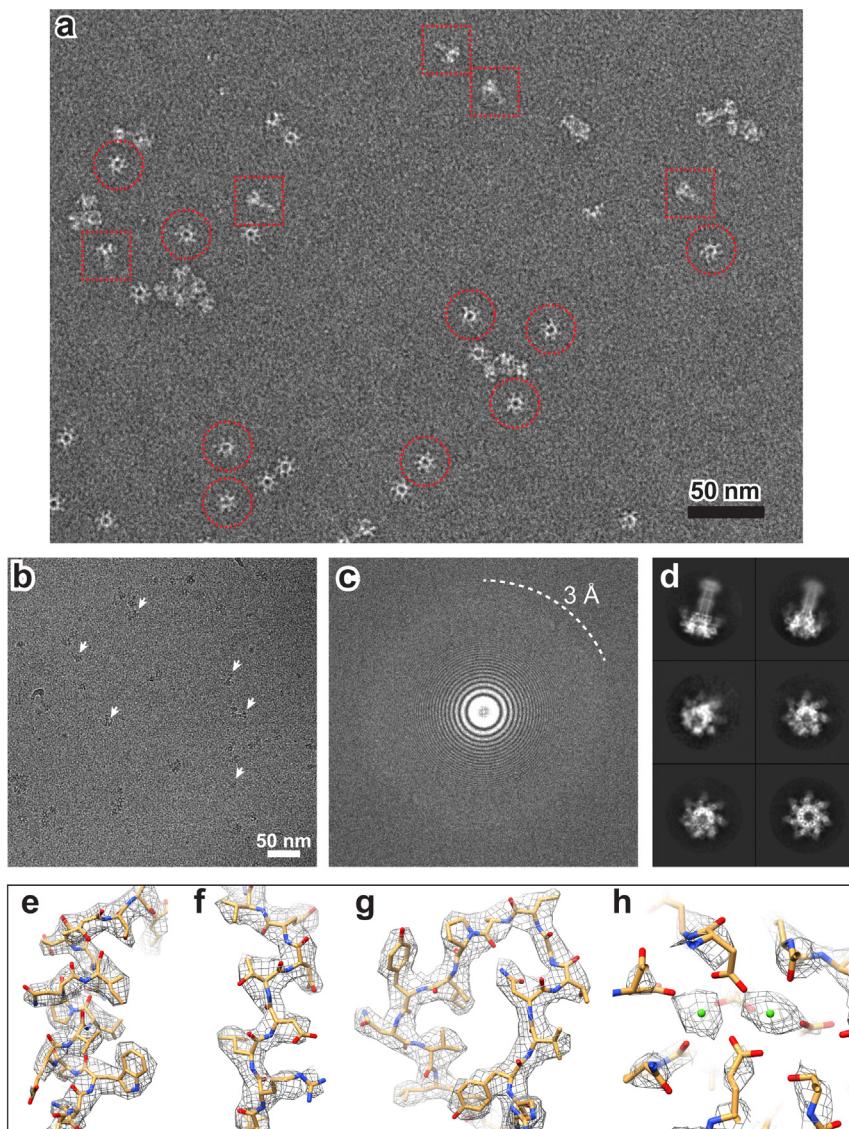
These 'shiny' particles were then subjected to three-dimensional auto-refinement in RELION to generate the final cryoEM map. RELION post-processing with a soft auto-mask⁴² estimated a resolution of 2.9 Å by the 'gold standard' Fourier shell correlation (FSC) at 0.143 criterion and a B-factor of -95 Å², and the post-processing without any mask reported a resolution of 3.3 Å. The accuracies of rotation and translation reported by RELION three-dimensional auto-refinement were 1.67° and 1.01 Å. For visualization and atomic model building, the cryoEM map was sharpened and low-pass filtered by RELION post-processing using the above-mentioned B-factor and resolution. Local resolution was calculated by ResMap⁴³ using the two cryoEM maps independently refined from halves of data.

Three-dimensional classification and auto-refinement also identified a subset of 21,632 particles that led to a cryoEM map at 3.6-Å resolution, which was in the pore state but lacked the density corresponding to the 14-stranded β-barrel: that is, it only had the corolla and calyx of the PA pore.

Atomic model building and refinement. *De novo* atomic model building of the PA pore except its domain 4 was performed on the cryoEM map at 2.9-Å resolution using Coot⁴⁴. The coarse model was then refined using PHENIX in a pseudo-crystallographic manner⁴⁵. Note this procedure only improved the atomic model and did not modify the cryoEM map. Briefly, the cryoEM map was put into an artificial crystal lattice to calculate its structure factor using the em_map_to_hkl.inp utility program in CNS⁴⁶. The amplitudes and phases of the structure factor were used as pseudo-experimental diffraction data for model refinement in PHENIX. The restraints of Ramachandran, secondary structure, and non-crystallographic symmetry were used in the refinement.

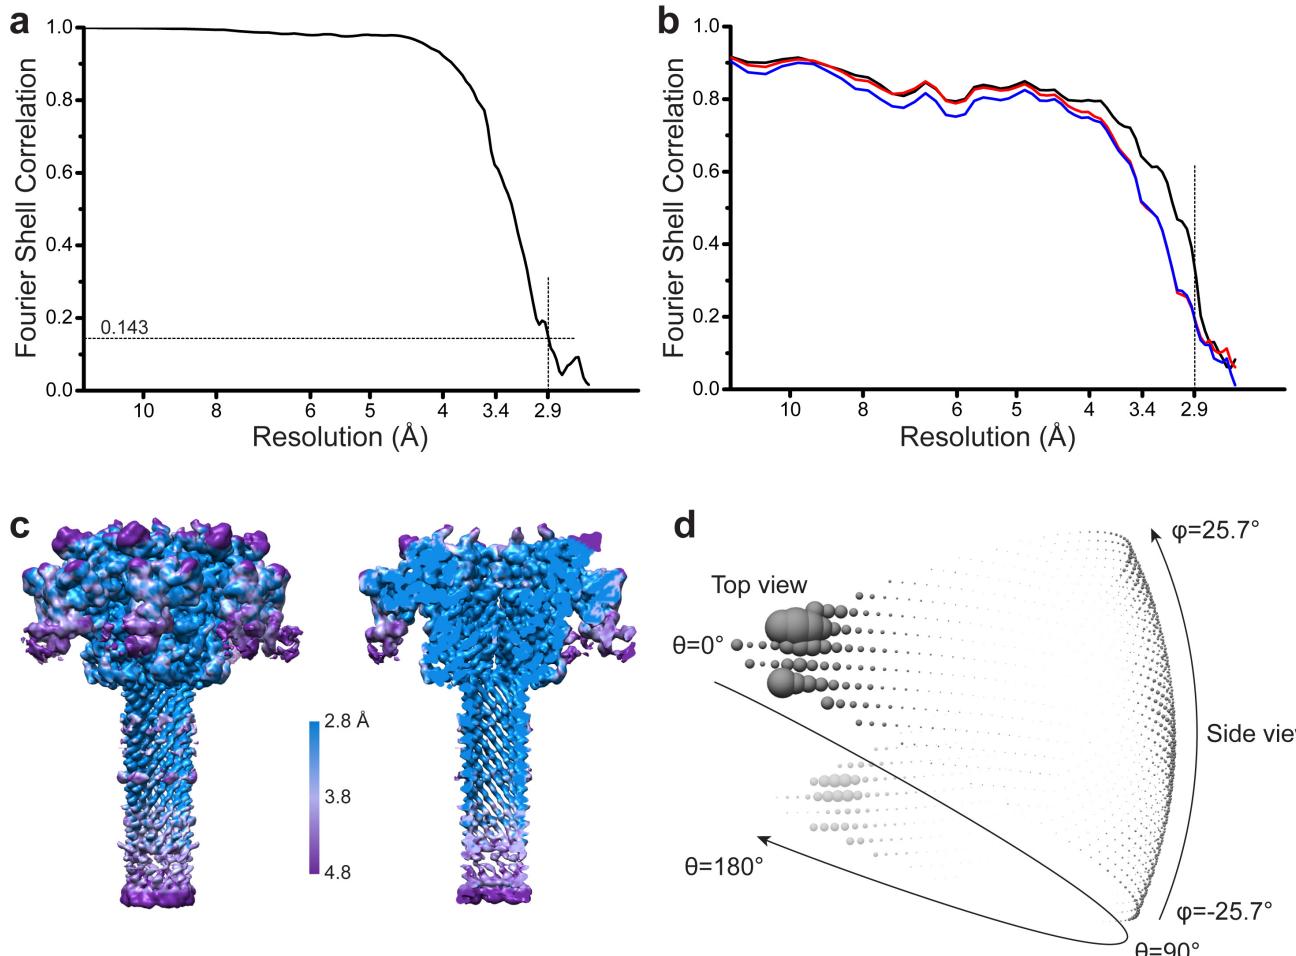
The cryoEM maps and atomic models were visualized using UCSF Chimera⁴⁷ or PyMOL⁴⁸. The available crystal structures of PA₈₃ monomer (for example, PDB accession numbers 1ACC and 3TEW) and PA₆₃ heptameric prepore (PDB accession number 1TZQ) are highly similar to each other; therefore 3TEW, which has the highest resolution among them, was used for structural comparison with our cryoEM structure of the PA pore in most situations and 1TZQ was used when the inter-protomer interaction or domain movement were considered.

31. Wigelsworth, D. J. *et al.* Binding stoichiometry and kinetics of the interaction of a human anthrax toxin receptor, CMG2, with protective antigen. *J. Biol. Chem.* **279**, 23349–23356 (2004).
32. Ortega, J., Singh, S. K., Ishikawa, T., Maurizi, M. R. & Steven, A. C. Visualization of substrate binding and translocation by the ATP-dependent protease, ClpXP. *Mol. Cell* **6**, 1515–1521 (2000).
33. Suloway, C. *et al.* Automated molecular microscopy: the new Leginon system. *J. Struct. Biol.* **151**, 41–60 (2005).
34. Suloway, C. *et al.* Fully automated, sequential tilt-series acquisition with Leginon. *J. Struct. Biol.* **167**, 11–18 (2009).
35. Li, X. *et al.* Electron counting and beam-induced motion correction enable near-atomic-resolution single-particle cryo-EM. *Nature Methods* **10**, 584–590 (2013).
36. Ludtke, S. J., Baldwin, P. R. & Chiu, W. EMAN: semiautomated software for high-resolution single-particle reconstructions. *J. Struct. Biol.* **128**, 82–97 (1999).
37. Mindell, J. A. & Grigorieff, N. Accurate determination of local defocus and specimen tilt in electron microscopy. *J. Struct. Biol.* **142**, 334–347 (2003).
38. Heymann, J. B. Bsoft: image and molecular processing in electron microscopy. *J. Struct. Biol.* **133**, 156–169 (2001).
39. Scheres, S. H. RELION: implementation of a Bayesian approach to cryo-EM structure determination. *J. Struct. Biol.* **180**, 519–530 (2012).
40. Scheres, S. H. A Bayesian view on cryo-EM structure determination. *J. Mol. Biol.* **415**, 406–418 (2012).
41. Scheres, S. H. Beam-induced motion correction for sub-megadalton cryo-EM particles. *eLife* **3**, e03665 (2014).
42. Amunts, A. *et al.* Structure of the yeast mitochondrial large ribosomal subunit. *Science* **343**, 1485–1489 (2014).
43. Swint-Kruse, L. & Brown, C. S. Resmap: automated representation of macromolecular interfaces as two-dimensional networks. *Bioinformatics* **21**, 3327–3328 (2005).
44. Emsley, P. & Cowtan, K. Coot: model-building tools for molecular graphics. *Acta Crystallogr. D* **60**, 2126–2132 (2004).
45. Adams, P. D. *et al.* PHENIX: a comprehensive Python-based system for macromolecular structure solution. *Acta Crystallogr. D* **66**, 213–221 (2010).
46. Brunger, A. T. *et al.* Crystallography & NMR system: a new software suite for macromolecular structure determination. *Acta Crystallogr. D* **54**, 905–921 (1998).
47. Pettersen, E. F. *et al.* UCSF Chimera—a visualization system for exploratory research and analysis. *J. Comput. Chem.* **25**, 1605–1612 (2004).
48. Schrodinger, LLC. The PyMOL Molecular Graphics System, v1.7.2 (2014).
49. Krissinel, E. & Henrick, K. Inference of macromolecular assemblies from crystalline state. *J. Mol. Biol.* **372**, 774–797 (2007).


Extended Data Figure 1 | Negative-stain and cryoEM of the PA pore.

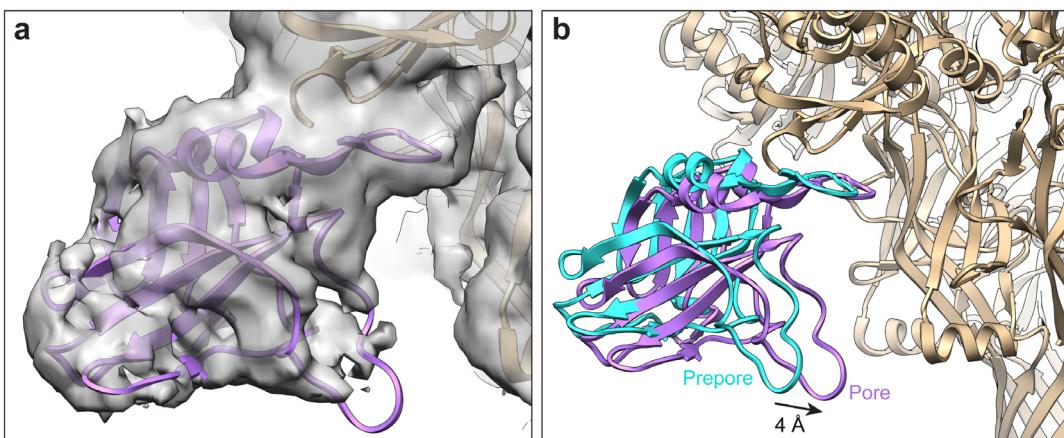
a, Negative-stain EM micrograph of PA pore particles. Some representative top-view and side-view particles are selected with circles and squares, respectively. **b**, A full-size drift-corrected cryoEM micrograph of 3,710 pixels \times 3,710 pixels of PA pore particles acquired from a Gatan K2 Summit direct electron detection camera, at 300 kV accelerating voltage, $-2.7 \mu\text{m}$ defocus, and

a total dose of 39 electrons per square ångström. Some representative side-view particles of PA pore are indicated by arrows. **c**, Power spectrum of the cryoEM micrograph in **b**. **d**, Representative cryoEM two-dimensional class averages of particles at different orientations. **e–h**, Superimposition of representative regions of the cryoEM map (mesh) with the atomic model (stick), including α -helix (**e**), β -strand (**f**), loop (**g**), and Ca^{2+} ions (green spheres in **h**).



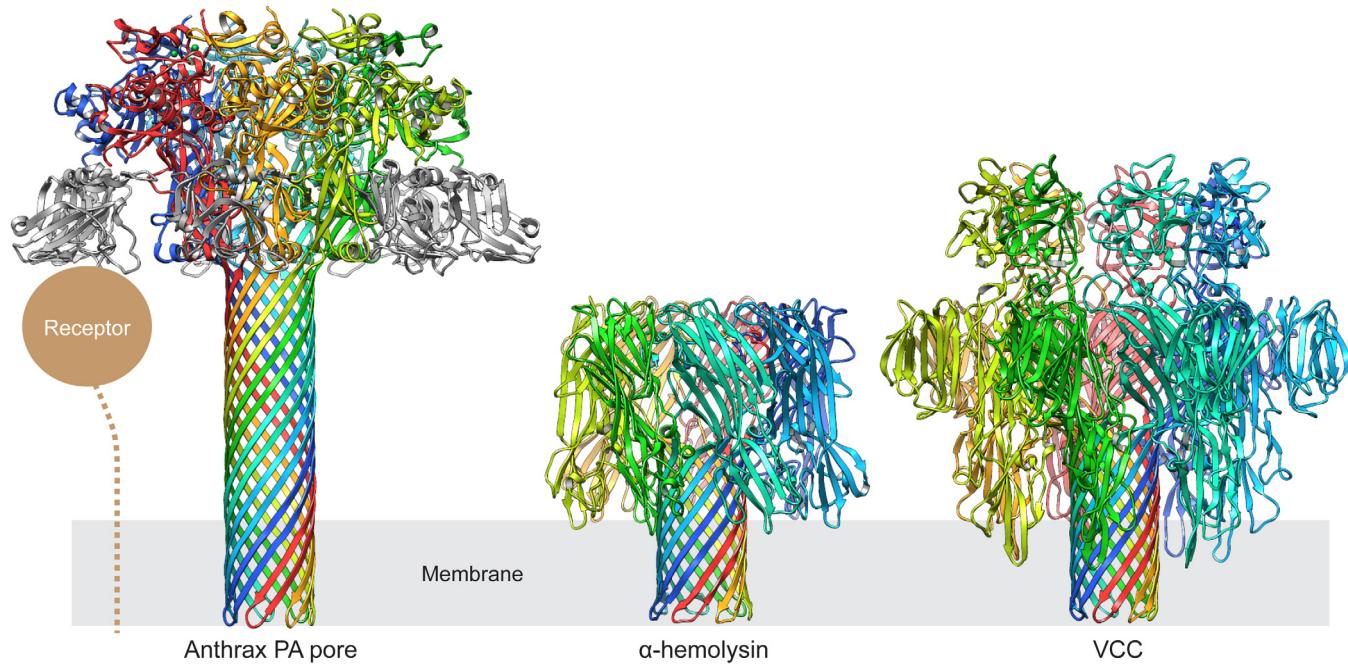
Extended Data Figure 2 | Resolution estimation of the cryoEM single particle reconstruction of the PA pore. **a**, ‘Gold standard’ FSC between two independently refined maps with an auto-mask that was corrected by phase randomization. The resolution was estimated by the ‘gold standard’ FSC at 0.143 criterion⁴². **b**, FSC of the final atomic model versus the final cryoEM map (black); of a model refined in the first of the two independent maps used for the

‘gold standard’ FSC versus the same map (red) and versus the second independent map (blue). **c**, Surface view and cut-through of the unsharpened cryoEM map coloured by local resolution estimated by ResMap⁴³. **d**, Euler angle distribution of all particles used for the final three-dimensional reconstruction.



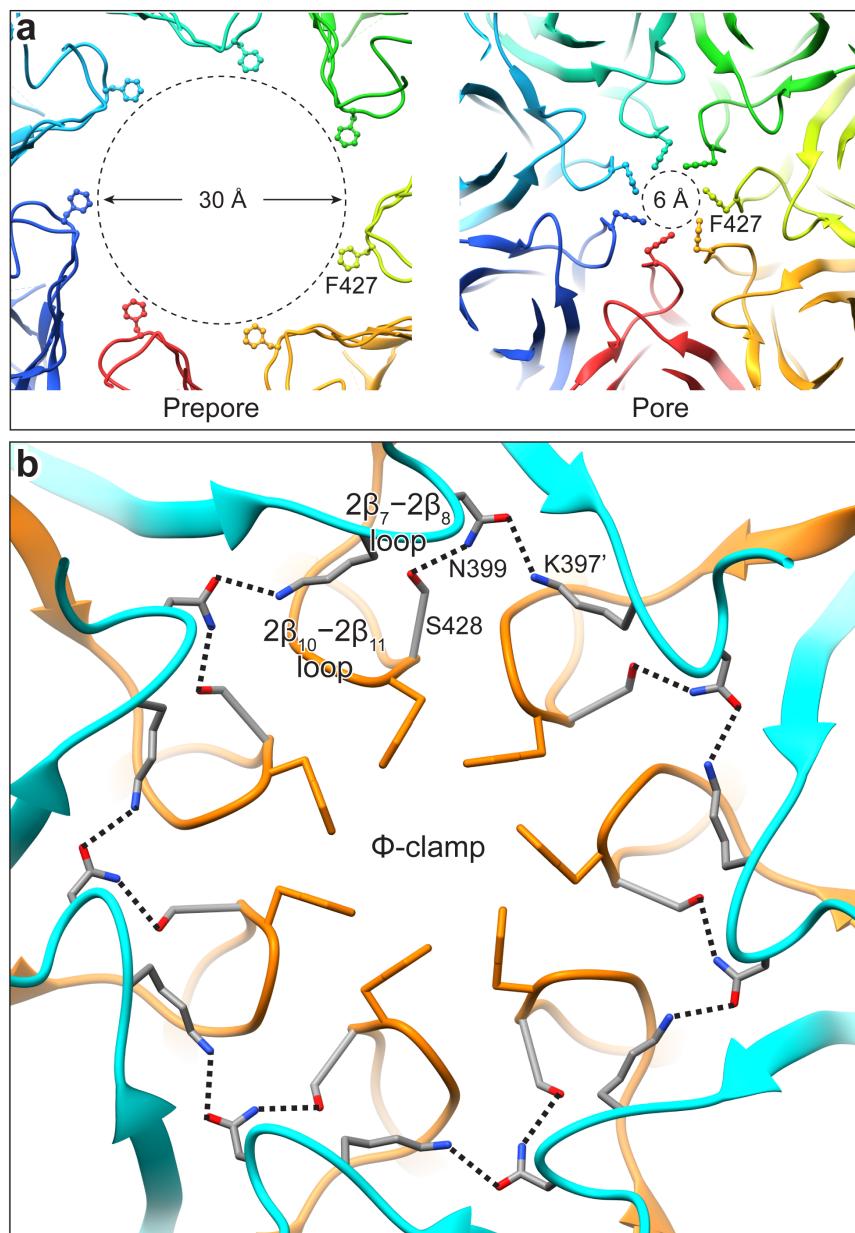
Extended Data Figure 3 | Fitting of domain 4 into the cryoEM map of PA pore. **a**, The crystal structure of domain 4 (purple ribbons) from PA prepore (PDB accession number 1TZO) fits the cryoEM map (grey surface) of PA pore with a good agreement (cross-correlation coefficient: 0.91).

b, Translocation of domain 4 during the prepore to pore conversion. Domain 4 is shifted inwards as a rigid-body for 4 Å from the prepore conformation (dark cyan) to the pore conformation (purple).



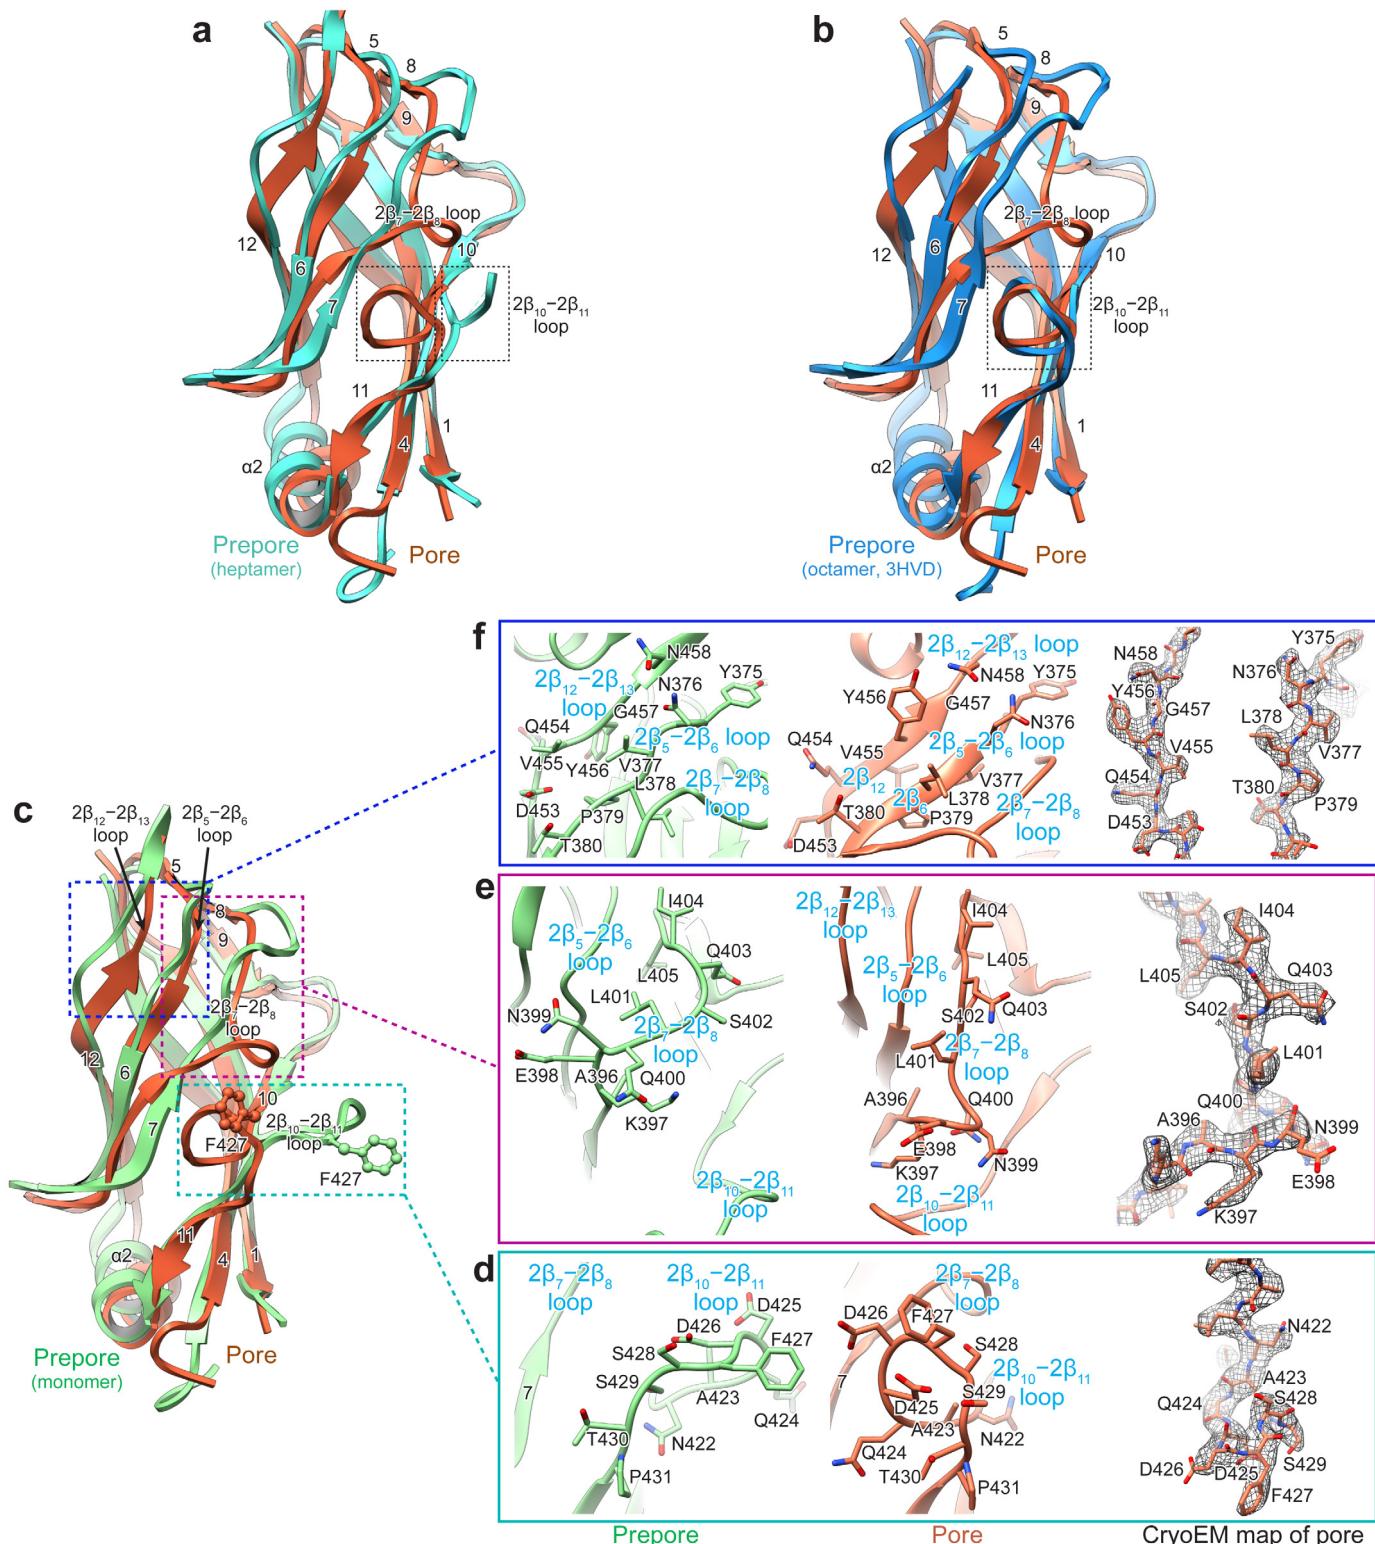
Extended Data Figure 4 | Comparison of the structures of the PA pore, α -haemolysin (PDB accession number 7AHL), and *Vibrio cholerae* cytolysin (PDB accession number 3O44). These three complexes are in their

membrane-inserted forms. The approximate position of the lipid bilayer is illustrated with grey shade. The host cell receptor (TEM8 or CMG2) of PA that binds to domain 4 (grey ribbons) is schematically illustrated.


Extended Data Figure 5 | Assembling of the Φ -clamp in the PA pore.

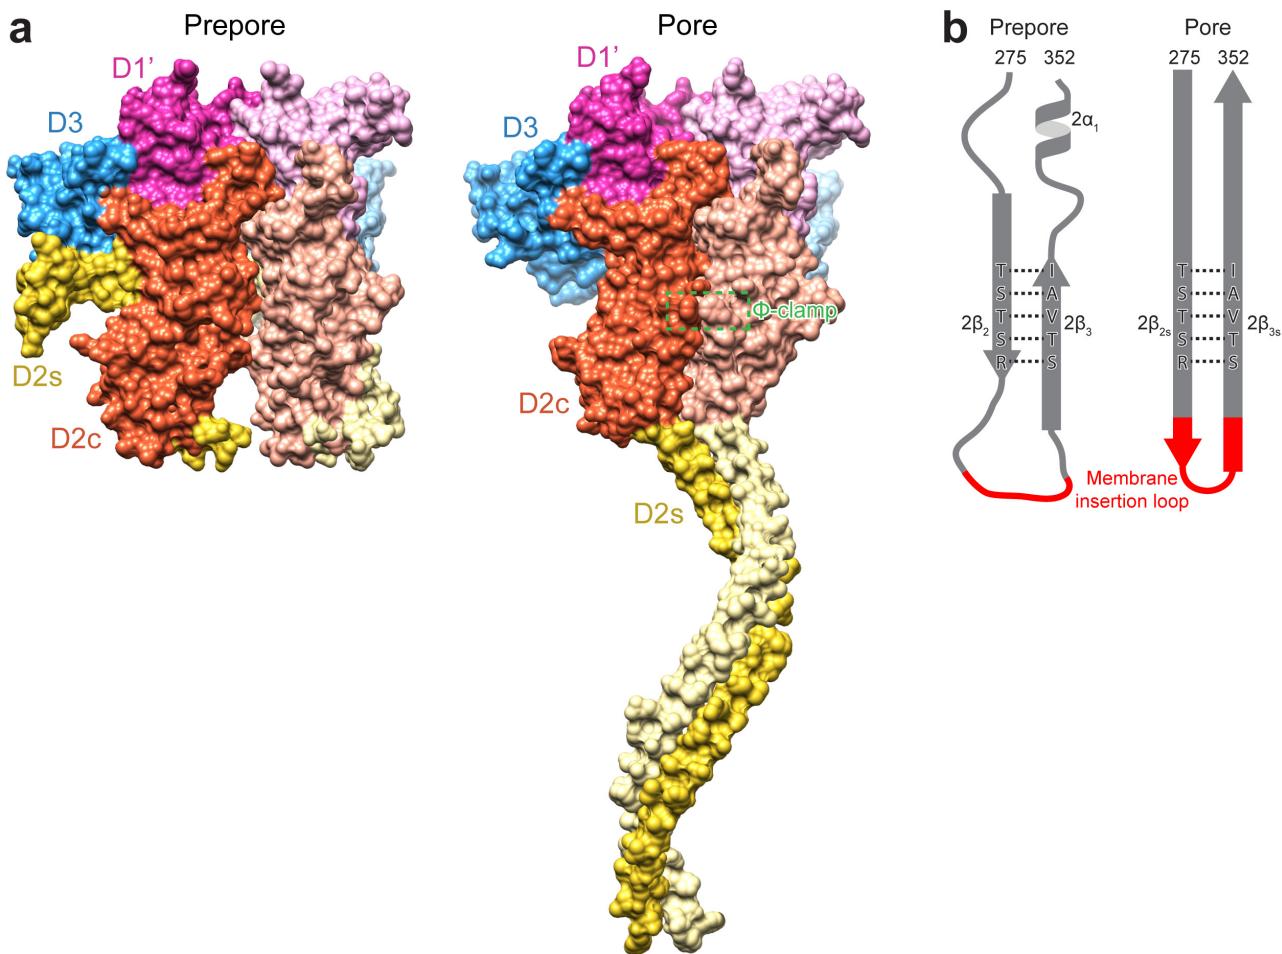
a, Comparison of the Φ -clamp in the PA pore with its corresponding region in the PA prepor. The residues Asp426 and Phe427 missing in the crystal structure of the PA₆₃ prepor (PDB accession number 1TZO) are modelled based the crystal structure of the PA₈₃ monomer (PDB accession number

1ACC). b, Close-up top view of the Φ -clamp region in the PA pore. The 2 β_7 -2 β_8 and 2 β_{10} -2 β_{11} loops are coloured in cyan and orange, respectively. The hydrogen bonds between Ser428, Asn399, and Lys397', which form a chain tethering the 2 β_7 -2 β_8 and 2 β_{10} -2 β_{11} loops together, are depicted with dashed lines.



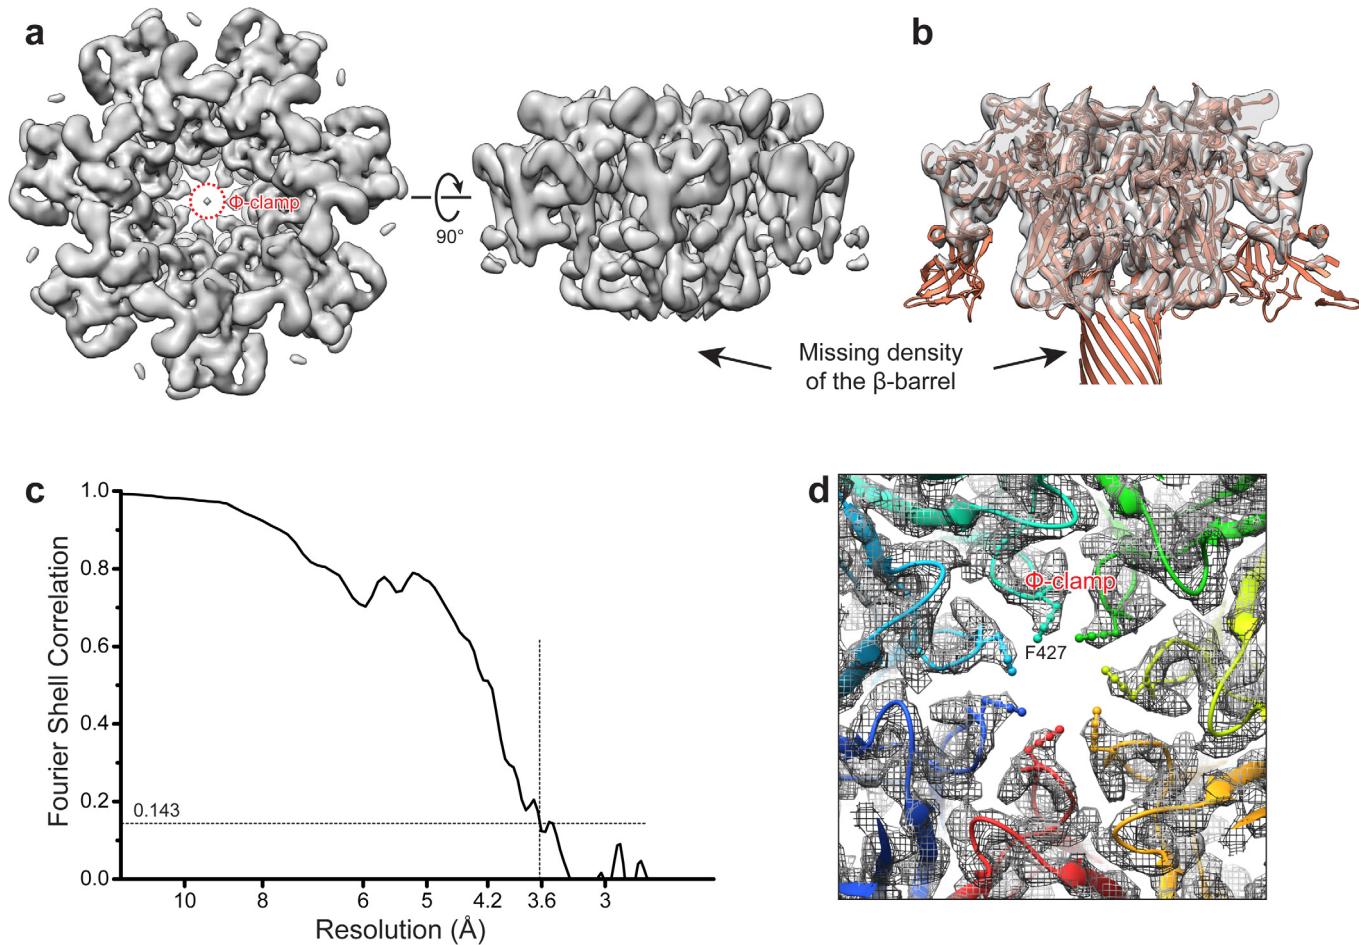
Extended Data Figure 6 | Conformational changes of domain 2c between the PA prepose and the PA pore. **a**, Superimposition of domains 2c of the PA pore and the crystal structure of the PA heptameric prepose (PDB accession number 1TZO). The dashed boxes highlight the conformational difference of the $2\beta_{10}-2\beta_{11}$ loop between these two structures. Residues Asp426 and Phe427 are not solved in the crystal structure of PA prepose because of their flexibility. **b**, Superimposition of domains 2c of the PA pore and the PA octameric prepose (PDB accession number 3HVD). Note the similarity of the $2\beta_{10}-2\beta_{11}$ loop (dashed box) between these two structures. **c**, Superimposition of domains 2c of the PA pore and the PA prepose (monomer; PDB accession number 3TEW). **d-f**, Close-up views of the

rearrangements of the $2\beta_{10}-2\beta_{11}$ loop (**d**), the $2\beta_7-2\beta_8$ loop (**e**), and the $2\beta_5-2\beta_6$ and $2\beta_{12}-2\beta_{13}$ loops (**f**) during the conversion from prepose to pore. Val377 and Leu378 of the $2\beta_5-2\beta_6$ loop and Val455 and Tyr456 of the $2\beta_{12}-2\beta_{13}$ loop are flipped upside down to obtain the β -strand conformation in the PA pore, leading to an extension of $2\beta_6$ and $2\beta_{12}$ (**f**). CryoEM densities (mesh) corresponding to these loops are displayed to the right of **d-f** with atomic models superimposed, showing unambiguous atomic modelling. The superimpositions of the cryoEM map and the atomic model are shown in views different from the left panels for clarity. The conformational changes of these loops result in a more compact domain 2c with a decrease of 567 \AA^2 in its surface area.



Extended Data Figure 7 | Compact assembly of the PA pore. **a**, Surface views of two neighbouring protomers (the same domains in the two protomers are in different shades of the same colours) of the PA pore and the PA prepor (PDB accession number 1TZ0) visualized from inside the PA heptamer. Domains 1', 2c, 2s, and 3 are coloured differently and domain 4 is not shown. The inter-protomer interface in the PA prepor is largely formed by domains 1' and 3, and domain 2 only contributes to this interface by its membrane insertion loop and carboxy (C)-terminal region ($2\beta_{13}$, $2\beta_{14}$, and $2\alpha_3$). In the PA pore, the convergence of domain 2c creates an inter-protomer interface without any gap, with an increase of interface area on domain 2c from $1,247 \text{ \AA}^2$ of the PA prepor to $2,106 \text{ \AA}^2$ of the PA pore as calculated using PISA⁴⁹ (<http://www.ebi.ac.uk/pdbe/pisa/>). Additionally, formation of the β -barrel also leads to extensive contacts, creating a new interface area of $1,195 \text{ \AA}^2$ between two protomers. **b**, Schematic of the conversion of $2\beta_2$, $2\beta_3$, $2\alpha_1$, and the membrane

insertion loop of the PA prepor to $2\beta_{2s}$ and $2\beta_{3s}$ of the PA pore. Hydrogen bonds between $2\beta_2$ and $2\beta_3$, which are depicted with dashed lines, are maintained during the conversion. By contrast, $2\alpha_1$ and the connecting loops have to be fully unfolded and converted into β -strands that collectively assemble the β -barrel of the PA pore. Although the detailed events of β -barrel formation are not yet clear, it is likely that assembly starts from the top in a zipper-like manner. A favourable scenario is that the convergence of domain 2c would place the top ends of the β -strands close to each other to form a short β -barrel, which could extend by pulling more residues together via formation of ordered hydrogen bonds until it reaches the bottom end. It is less favourable that the assembly starts from other regions because disordered hydrogen bonds and hydrophobic interactions could generate enormous non-productive, possibly irreversible pairings between strands.



Extended Data Figure 8 | CryoEM single-particle reconstruction of a subset of particles lacking the 14-stranded β -barrel. **a**, Surface views of the cryoEM map lacking the β -barrel. The map is unsharpened and has a resolution of 3.6 Å. **b**, Cross-section side view of superimposition of the unsharpened cryoEM map with the atomic model of the intact PA pore, showing the cryoEM map has the same conformation as the PA pore except for the absence of the 14-stranded β -barrel. **c**, ‘Gold standard’ FSC (with an auto-mask that was corrected by phase randomization) between two independently refined maps.

d, Top view of the Φ -clamp region of the cryoEM map (mesh) lacking the β -barrel superimposed with the atomic model of the PA pore (ribbons), showing correct assembling of the Φ -clamp in the cryoEM map.

Extended Data Table 1 | Data collection and structure refinement statistics

Data Collection	
Particles	60,455
Pixel size (Å)	1.28
Defocus range (μm)	-1.8 to -5.1
Voltage (kV)	300
Total electron dose	39 e ⁻ /Å ²
Electron dose used in the final reconstruction	30 e ⁻ /Å ²
Refinement	
Resolution (Å)	2.9
Map sharpening B-factor (Å ²)	-95
Average B-factor (Å ²)	81
R.m.s deviations	
Bond lengths (Å)	0.009
Bond angles (°)	1.083

CAREERS

POSTDOCS A lab head looks for trainees, not employees p.553

NATUREJOBS IN BOSTON Entrepreneurship and self-direction go.nature.com/wyubyi

 **NATUREJOBS** For the latest career listings and advice www.naturejobs.com

DANIEL STOLLE



MEDICAL RESEARCH

Subject to reflection

Most scientists who study disease carry out their research with an eye to treating others — but a few have only to look at their own bodies to feel the need for a cure.

BY VIRGINIA GEWIN

Lukas Wartman always wanted to be a physician-scientist — but he never expected to be a research subject. Nor did he anticipate that his postdoc adviser would end up sifting through his genome for clues to treatment.

In 2003, during his final year of medical school, Wartman was diagnosed with the blood cancer acute lymphoblastic leukaemia (ALL). He went into remission after chemotherapy, relapsed one year into his clinical fellowship in oncology and then recovered following a bone-marrow transplant. In 2008, he started postdoctoral research in cancer genetics at Washington University in St Louis (WUSTL), Missouri.

At the time, his adviser, cancer researcher Timothy Ley, was carrying out the first whole-genome sequencing study of acute myeloid leukaemia (AML). When Wartman relapsed again in 2011, Ley and his colleagues sequenced his cancer genome, too. They found that a certain gene was overactive, which pointed them to sunitinib, a drug for kidney cancer that is known to reduce the gene's activity.

Despite having to cope with complications from a second bone-marrow transplant, Wartman continues to conduct cancer research, emboldened by his own illness. He focuses on AML, which produces symptoms similar to those of his cancer, but has a different genetic cause. Last December, Wartman published preliminary data on the suitability of a drug for

both AML and ALL in the American Society of Hematology journal, *Blood* (L. Wartman *et al.* *Blood* **124**, 5292; 2014).

Scientists such as Wartman, who research diseases that affect themselves or their loved ones, occupy a curious niche in the scientific enterprise. Their experiences can offer unique research insight, garner media attention and provide valuable connections to patient groups. And they are highly motivated to find a cure.

But their jobs are also fraught with the emotional — and sometimes ethical — challenges that can arise when researching something that resonates deeply on a personal level. Several of Wartman's peers and mentors encouraged him to leave oncology to free himself from thinking about cancer all the time. ▶

► Those who come to know a disease as both patient and investigator can take steps to preserve their emotional health and research objectivity. These scientists need to be mindful of their motivations and realistic about their goals — and foster a broad and deep support network for times when research findings hit too close to home.

SOUL MEETS BODY

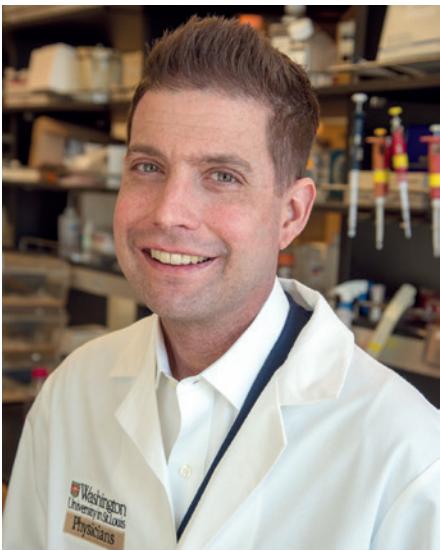
When a personal connection is involved, the drive to improve research into and treatment of a disease can be especially powerful. Geneticist Angela Christiano knows this first-hand: she jumped into her line of research after being diagnosed with an autoimmune disease called alopecia areata, which causes hair to fall out. As a postdoc, she had worked on dermatological diseases, but she began to look for another research focus when she moved into a tenure-track position at Columbia University in New York City. One day in 1996, her hairdresser pointed out a bald spot on the back of her head. After her diagnosis, she learned that frustratingly little was known about her condition.

And so she began a decade-long journey with patient-advocacy groups to establish a registry that now contains 3,000 serum and cell samples for genetics work. In 2008, she conducted a genome-wide association study and found seven genes linked to the disease. She also uncovered similarities between alopecia areata and other autoimmune disorders such as type 1 diabetes and coeliac disease.

Christiano launched clinical trials to test whether two drugs — one for rheumatoid arthritis and another for a bone-marrow disorder — could also combat alopecia areata. Last year, she published results showing that three patients treated with one of these drugs achieved near-full hair regrowth (A. Christiano *et al. Nature Med.* **20**, 1043–1049; 2014). Her emotional bond with the patient community fuels her work. “I give talks every year to the National Alopecia Areata Foundation patient conference. The year that I presented the genome-wide association study — and when I announced drug candidates for clinical trial four years later — they gave me a standing ovation,” she says. “We were all sobbing.”

A deep personal drive might be just what is needed in an understudied, underfunded field. Leonard Jason, a psychologist who directs clinical training, found his diagnosis of chronic fatigue syndrome (CFS) in 1990 to be life-altering, both personally and professionally. After an 18-month leave of absence from DePaul University in Chicago, Illinois, to recuperate, he delved into the CFS literature and realized a need for a better definition of the illness, a less-stigmatizing name and more-robust research on prevalence.

The field was riddled with questions — indeed, many physicians dispute whether CFS is even a genuine illness. “A number of folks had their careers destroyed by coming into this



Lukas Wartman studies the genetics of cancer.

area,” Jason says. Cognizant of potential career fallout and his susceptibility to overexertion, he focused on what the field needed most and what he could realistically achieve. He spent a decade on epidemiological studies that expanded estimates of the US patient population from 20,000 to 900,000. His work also showed that medical interns are more dismissive of the term chronic fatigue syndrome than of the more physiologically based name, myalgic encephalomyelitis (L. Jason *et al. Am. J. Community Psychol.* **30**, 133–148; 2002). “The key is to keep the focus on small wins,” he says. “I’ve been able to work with patients, researchers and government officials for over two decades on a variety of topics — from changing the name, to searching for biological markers.”

Jason might be one of only a few people with CFS conducting research on the condition, but that is not the case in other areas. “In type 1 diabetes research, there is an overrepresentation of people with the disease,” says Timothy Tree, an immunobiologist at King’s College London. And Tree is one of them.

Tree is driven by one specific question — why did he, and not his brothers, inherit the disease from their father? His need to answer that question, as well as whether he could pass it on to his children, keeps him going in the lab. “When the experiment goes wrong or funding gets turned down, having the disease gives me a reason to stay in it for the long haul,” he says.

Researchers in his position can offer practical perspectives, Tree says. For example, if a colleague suggests an idea for a therapy, he can give an informed opinion on how likely patients would be to comply. “It gives me a different kind of objectivity.”

Yet some people are concerned that such personal significance could compromise objectivity by introducing bias. “Bias can operate at a subconscious level,” says David Resnik, a bioethicist at the US National Institute of Environmental Health Sciences in Research

Triangle Park, North Carolina. And the results can potentially skew the work, he adds. “It might lead a researcher to make interpretations of data or study designs differently.”

Rebecca Dresser, a research ethicist at WUSTL, says that researchers who study their own or family members’ diseases can take precautions to safeguard against bias or perceptions of bias. “Consult with people who don’t have the same personal perspective — and get them on the team reviewing grants and research protocols,” she says.

But researchers such as Christiano and Tree say that they can put scientific integrity ahead of hope. If anything, Christiano says, having alopecia areata made her more cautious and conservative. “I’m the most demanding one on the team. I want to be triple sure about any findings before we go forward.” And Tree notes that his experience with diabetes has taught him how important it is not to oversell positive results. “As a patient, I know that people hang on to every word you say for hope.”

A FINE LINE

Balancing work and personal life is always important — more so when that work can have an impact on a scientist’s emotional state as well as productivity. “Having a bit of distance — between my own problems and the larger-scale problems we were trying to understand in the lab — definitely made things easier for me,” says Wartman. Of course, he does not directly study his own genome. “Tim and I sat down and had a serious discussion about whether or not even doing the genomics of AML would be too psychologically heavy for me to take on,” he says.

Wartman’s experience has strongly influenced his priorities. He benefited from the

“As a patient, I know that people hang on to every word you say for hope.”

sequencing of his cancer genome, and feels a need to help others to gain similar insight. So he chose to devote his energy to the genomics tumour ward at WUSTL. Each month, he and his colleagues invite cancer-care professionals to present cases. Together, they decide which diseases merit genome sequencing, bridging the gap between research and discovery genomics. “I spend a lot of my off hours getting this up and running,” he says.

But sometimes the pressure can become too great. Michael Dodd was on track to conduct his postgraduate research at the University of Oxford, UK, on a condition that afflicts his father called hypertrophic cardiomyopathy, in which the heart muscle thickens. He did not know that he, too, carried the underlying genetic mutation, until a year into his PhD programme when a genetic test came back positive. “It was quite a shock to find out, especially since, to this day, I don’t present any symptoms,” he says. At one point, his supervisor, Hugh Watkins, doubled as his physician.

Dodd chose to shift his research focus elsewhere. "I sometimes found it weird to be in the lab," he says. He was one of several patients who had the mutation, yet no symptoms, and so had MRI scans in their lab. "It was weird to see a bar graph, knowing I'm one of the points," he says.

The research could be emotionally taxing. "It would feel odd to work on, for example, a mouse with the same genetic mutation as me, and wonder if I would respond similarly," he says. But he did want to keep working on the heart, so he is now a postdoc studying the cardiac effects of diabetes, a disease that his grandfather had.

SPOTLIGHT SCARS

The emotional toll can be especially intense when media attention forces the scientist into the public eye. Wartman felt the landscape shift after a high-profile piece about him appeared in the *New York Times* in 2012. He is happy that patients find his personal perspective helpful, but regrets that the decision to share his story no longer rests with him. "It's still not the easiest topic for me to talk about," he says. "The last time I relapsed, I came close to dying. To rehash that on a regular basis is emotionally draining."

Media attention can change one's entire research career. Kay Redfield Jamison, a clinical psychiatrist and founder of a clinic for mood disorders at the University of California, Los Angeles, channelled her struggles with bipolar disorder into research on the illness's wide range of effects — from enhanced creativity to a high risk of suicide. But when she wrote her autobiography in 1995, entitled *An Unquiet Mind: A Memoir of Moods and Madness*, she knew that her professional life would never be the same. She gave up her clinical practice. "You can't say that you've been psychotic and nearly died by suicide and expect people to look at you the same way," she says.

Now at Johns Hopkins University in Baltimore, Maryland, Jamison focuses on writing and public speaking. She credits a network of supportive friends and colleagues for helping her to navigate her career ups and downs. "Becoming a poster child for an illness is draining," she says. "It becomes a disturbing part of your identity." Still, it was worth it to reach others who were suffering. "That's what good comes out of it."

At the end of the day, that desire to aid others motivates many researchers to continue their work even though their own health is poor. "Leukaemia disrupted my career and goals and was a huge setback in my life," Wartman says. "At the same time, if I can turn my own struggle into a story that helps other people, that has value." ■

EINAT SEGEV

Virginia Gewin is a freelance writer in Portland, Oregon.

TURNING POINT

Roberto Kolter

Roberto Kolter set up his microbiology laboratory at Harvard Medical School in Boston, Massachusetts, in 1983. Postdocs worldwide hope to join his lab because of his career-targeted training philosophy, but with rare exceptions, he brings in only those who already have a fellowship.

Why do you accept postdocs only if they have their own funding?

I focus on those whom I believe have a fantastic chance of getting their own funding as a principal investigator. I think it's unfair for me to interview those who have very little chance of getting their own funding, considering how competitive the academic job market is and how important it is to show independence.

What does your laboratory focus on?

I let the postdocs explore what they want to explore, as long as it is within the sphere of my interest. I've worked on starvation physiology, biofilms, signalling, experimental evolution, antibiotics and many other subjects.

Describe your training philosophy.

I train people to go on into academia, industry, the corporate world or whatever they want to go into. We need to give them the experience that they require, including learning how to teach and learning how to manage. Postdocs are not just there to come to the lab so that principal investigators can get their next grant.

What stands out when you look at applications?

I have learned that networking works very, very well. If I know who trained that individual, and I know and respect them, then I'll know a lot about how this postdoc will work in the lab. But that does not mean that if I don't know the mentor I will close the door to the postdoc. They need to have also done their homework — they need to know how I train people and how they think they would fit in.

When have you made exceptions?

There are one or two cases where I was completely sure that they would get a fellowship, and they didn't. But by then I had gotten so excited about the project we had co-developed that I chose to support them from my own funds.

How does your lab develop a research project?

The ideas often emerge from conversations that start about 18 months before the postdoc comes to work with me. It has almost always been my policy that incoming postdocs build their research projects and are free to take the



project with them once they leave, to help them to set up their own lab. That gives the postdocs who are leaving a good opportunity to establish themselves without having to compete with me and the people in my lab.

What careers do your postdocs pursue?

About half the 100 or more postdocs that have gone through my lab hold full-time academic jobs, of which running a research lab is a big component. Many people whom I take on as postdocs want a job in the biotechnology arena. The other 50% are dominated by those who choose to join a company. Those can range from start-up biotech companies to very well-established pharmaceutical or chemical companies. Others lead research groups at institutes or government labs, work as research associates, teach science or do other science-related work. Only two have left science.

Do they get permanent positions right away?

No one who has come through my lab has had to leave science because they could not get a job. Personally, I believe that I have failed a postdoc if I take them into my lab and they cannot get a job that they love when they leave. That usually means that they have to go on to do a second postdoc. There have been very few such individuals — fewer than five, in the 32 years I have had my own lab. So overall I rate my success rate in helping postdocs get their first job at about 90%.

What do you see as the role of a postdoc?

The meaning of postdoctoral training has been lost in today's scientific community. As mentors, we need to really reconsider what we are training postdocs for. And that's just it: it's a training period, not a job. ■

INTERVIEW BY JULIE GOULD

AN EXCERPT FROM *DYING FOR DUMMIES (2020)*

A guide to the next stage of your life.

BY NORMAN SPINRAD

INTRODUCTION

It used to be that choosing how to die involved no complicated menu of alternatives. You waited for it to happen or you took your own life — and that was it. And if you didn't want to decide on how to dispose of the material remains, you didn't have to — you could just leave that choice to your next of kin. But scientific advances have created more interesting and much more positive post-mortem modes to choose from. That's the good news.

The bad news is that making the choice requires sophisticated technological knowledge that is beyond what the average person is likely to have. You must understand that the choice is, in effect, the most crucial one you have to make, not merely in your life but potentially for ever. And you need to realize that the consequences will not be assured, that post-mortem existence is not a matter of certainty, but a bet on the odds that each mode currently seems to offer.

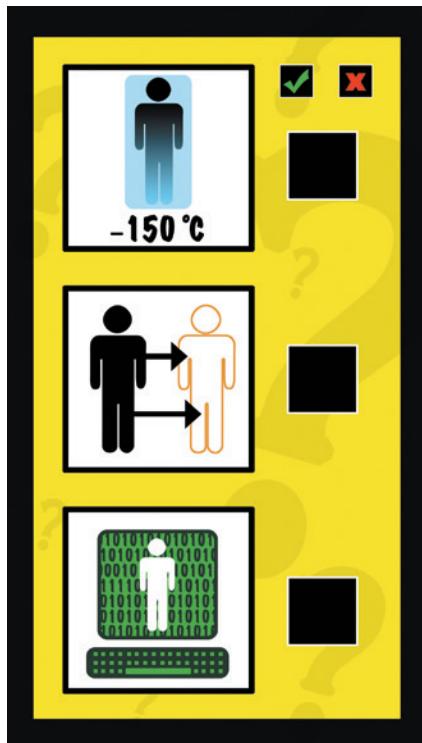
So the purpose of *Dying For Dummies* is to allow the public at large to make this choice in an enlightened and informed manner by explaining those odds and the costs of each post-mortem mode of existence.

CHAPTER ONE: THE CURRENT POSSIBILITIES

There are currently three post-mortem modes to choose from:

1) Cryogenic freezing of either the whole body or just the head. This is the most expensive choice, because the cryogenic state must be maintained indefinitely until the technology to revive your body or at least your brain is developed, if ever. In addition, it would be best to be frozen before you are legally dead, but that is not legal in all jurisdictions. So, given the expense and the odds, this cannot be recommended for most people.

2) Creation of a cloned body with your own genome and a brain transfer from your original body into it. This surgery is already possible and is favoured by those who believe that consciousness or the so-called 'soul' is an artefact of the unique individual cerebral meatware. Although initially more expensive than cryogenic freezing, it is cheaper in the medium run but, as with any major operation,



not without significant risk — and in the long run it is not a permanent solution, as the cloned body will eventually wear out too, and the procedure will have to be repeated.

Also, the choice of retaining the original cerebral meatware would seem to mean that the brain will continue to deteriorate in the cloned body or sequence of cloned bodies, although this has yet to be proved.

Still, this would be the obvious choice for those with religious or philosophical belief systems contending that if the 'soul', i.e. the consciousness software, is not running on the original meatware, 'you' are not really there.

3) Upload the instantaneous hologram of your consciousness and your memory data at or before the moment of death into a material matrix more robust than protoplasm — that is, computer hardware. Currently this seems to work to the point

at which the 'entity' in the computer can easily enough pass a 'Familial Turing Test'. Next of kin or lover

of the materially deceased speaks with the uploaded entity in the computer and with a synthesized emulation without being told which is which and correctly identifies the true post-mortem consciousness about 70% of the time, significantly above random.

However, the technology to download the consciousness software in the computer to a cloned brain in a cloned body has not yet been perfected beyond experiments with simplified backups and with equivocal results. Some of the software entities in the computer hardware declare that they don't even want to be downloaded into disgusting squishy protoplasm and are more than content to wait until metal bodies become available. And some of them don't even want those 'bodies' to emulate protoplasmic human bodies, and instead are eager to become aeroplanes, spaceships, submarines and even more outré customized 'bodies'.

This choice would seem to be the best one for those of modest means, as the upload to computer memory is the cheapest option, but it is probably only for the adventurous and those who believe that their essential being is the software of their consciousness independent of the meatware or hardware it is running on.

However, if consciousness is data and software independent of the meatware or hardware it is running on, it can be duplicated and backed up. But if it can be duplicated once, it can be duplicated many times and, once the technology is perfected, downloaded into multiple bodies, metallic creations, cloned protoplasm, or all of the above.

The practical advantages would be obvious. 'You' can be all the whatevers you want to be and at the same time. 'You' can be a whole football team, or a whole board of directors, or a whole symphony orchestra, or the whole population and machineries of an orbital habitat or otherwise generational starship.

So the question is:

CHAPTER TWO: WHICH IS 'YOU'?

... to be continued. ■

Norman Spinrad has been publishing novels in English for an actual half-century. His latest publication in English is the pamphlet *Raising Hell*. His latest novel has just been published in French as *Police Du Peuple*.

ILLUSTRATION BY JACEY